University of Groningen

Numerical continuation of normally hyperbolic invariant manifolds

Broer, H.W.; Hagen, A.; Vegter, G.

Published in:
Nonlinearity

DOI:
[10.1088/0951-7715/20/6/011](https://doi.org/10.1088/0951-7715/20/6/011)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2007

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Broer, H. W., Hagen, A., & Vegter, G. (2007). Numerical continuation of normally hyperbolic invariant manifolds. *Nonlinearity*, 20(6), 1499-1534. <https://doi.org/10.1088/0951-7715/20/6/011>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Numerical continuation of normally hyperbolic invariant manifolds

H W Broer¹, A Hagen^{2,3} and G Vegter¹

¹ Department of Mathematics and Computing Science, University of Groningen, PO Box 800, 9700 AV Groningen, The Netherlands

² Department of Mathematics, University of Texas at Arlington, 411 S. Nedderman Dr., Arlington, TX 76019, USA

E-mail: hagen@nethere.com

Received 6 October 2006, in final form 13 April 2007

Published 4 May 2007

Online at stacks.iop.org/Non/20/1499

Recommended by A Chenciner

Abstract

This paper deals with the numerical continuation of invariant manifolds regardless of the restricted dynamics. Common examples of such manifolds include limit sets, codimension 1 manifolds separating basins of attraction (separatrices), stable/unstable/centre manifolds, nested hierarchies of attracting manifolds in dissipative systems and manifolds appearing in bifurcations. The approach is based on the general principle of normal hyperbolicity, where the graph transform leads to the numerical algorithms. This gives a highly multiple purpose method. The graph transform and linear graph transform compute the perturbed manifold with its hyperbolic splitting. To globally discretize manifolds, a discrete tubular neighbourhood is used, induced by a transverse bundle composed of discrete stable and unstable bundles. This approach allows the development of the discrete graph transform/linear graph transform analogous to the usual smooth case. Convergence results are given. The discrete vector bundle construction and associated local k -plane interpolation may be of independent interest. A practical numerical implementation for solving the global equations underlying the graph transform is proposed. Relevant numerical techniques are discussed and computational tests included. An additional application is the computation of the ‘slow-transient’ surface of an enzyme reaction.

Mathematics Subject Classification: 34C, 37D, 37M, 41A, 65D, 65L, 65P, 70K

³ Current address: 1255 Van Buren Avenue, St. Paul, MN 55104, USA.

1. Introduction

Invariant manifolds give information about the global structure of phase space. For example, a codimension 1 manifold may separate several basins of attraction. Invariant manifolds are also used to simplify dynamical systems. The phase portrait near the manifold is trivial, so restricting the dynamical system to the manifold effectively reduces the dimension of the system. In some cases invariant manifolds can give a complete qualitative description of phase space. An example is a nested hierarchy of attracting manifolds in a dissipative system whose global attractor is a fixed point. If the global attractor is more complicated, it may be contained in an attracting manifold which contains the non-trivial dynamics. A simple example of the kind of thing we are after is a 3-torus attractor with phase lock dynamics, where the aim is to visualize an unstable 2-torus separatrix inside this 3-torus. In the ambient phase space, both the 3-torus attractor and the 2-torus, which is of the saddle type, may be computed using the method of the present paper.

The key notion needed here is *normal hyperbolicity* of the invariant manifold. According to the invariant manifold theorem, this guarantees the smooth persistence of the manifold under small perturbations of the system. Normal hyperbolicity generalizes the linearization method for hyperbolic fixed points and periodic orbits to higher dimensional manifolds. For an invariant submanifold V of a Riemannian manifold M , the normal hyperbolicity of V is exhibited by a splitting of $T_V(M)$ into invariant stable, unstable and centre parts. The persistence property of normally hyperbolic manifolds enables the development of insensitive numerical algorithms that compute the manifolds by numerical continuation.

1.1. What is computed?

In the present paper, the focus is on a method for computing (non-local) invariant manifolds of dimension ≥ 2 . Some of the related work in this category concerns quasiperiodic (for example, [22]) or attracting (for example, [12]) tori, parts of global attractors [10] or global (un)stable manifolds [28]. See also [29]. The computations of tori use global parametrizations of the tori where simplicial complexes are used in the present paper. The computations of parts of global attractors use successive subdivisions of a covering of a part of the global attractor. Starting from a given bounded region, this approach computes global attractors which are smooth or non-smooth. The computations of global (un)stable manifolds are concerned with extending a given piece of the manifold, to fill out the global (un)stable manifold.

This paper is concerned with methods which fit into a continuation context. These compute an invariant manifold of a perturbed system from initial data given by a manifold for the unperturbed system [2–4, 12, 19, 22]. This type of approach is motivated by the classical perturbation theory of invariant manifolds, as embodied in the invariant manifold theorem. Given a diffeomorphism F and an F -invariant submanifold V , the goal is to compute the invariant manifold $\tilde{V}$ for a nearby diffeomorphism $\tilde{F}$.

1.2. How is it computed?

An algorithm for computing an invariant manifold is based on one of the two types of invariance condition. This is because an invariant manifold V of a dynamical system is characterized in one of two ways. The vector field invariance condition states that at every point p in V , the vector field is tangent to V at p . The diffeomorphism invariance condition states that the image of V under the diffeomorphism is equal to V . An invariant manifold for a vector field may be characterized by the diffeomorphism condition using the time- T

map of the flow, while an invariant manifold for a diffeomorphism may be characterized by the vector field condition using the suspension of the diffeomorphism [43]. The classical approaches to proving the general existence of normally hyperbolic invariant manifolds are the graph transform [14, 27] or the Lyapunov–Perron mapping [20, 21], which use the diffeomorphism condition, and the PDE approach (successive approximation by solutions of linear PDEs) [44], which uses the vector field condition. The vector field invariance condition is a first order PDE, while the diffeomorphism condition uses the integrated vector field. Thus, the invariance condition used is a primary factor influencing the character of the resulting algorithm.

1.2.1. General approaches. In current computer algorithms, the diffeomorphism invariance condition for the perturbed manifold is usually solved using the graph transform or a functional Newton’s method. The vector field invariance condition for the perturbed manifold is typically solved using a functional Newton’s method. These three approaches are outlined next.

The graph transform obtains successive approximations to an invariant manifold for a perturbed system $\tilde{F}$ as follows. The basic idea is to split the system in a way which respects the dynamics. Given a normally hyperbolic F -invariant manifold V and a transverse bundle $N(V)$, the $\tilde{F}$ dynamics are split into a part in V and a part transverse to V . Given a section σ of $N(V)$, the graph of σ is transformed by the $\tilde{F}$ dynamics, obtaining a section σ' giving a closer approximation to the perturbed manifold $\tilde{V}$. For example, in the attracting case, the section σ' is obtained by solving two coupled equations which split the $\tilde{F}$ -invariance condition into a part in V (the *restricted equation*) and a part transverse to V (the *transverse equation*). The main part is to solve the restricted equation for a globally unique solution $q \in V$, for each $p \in V$. Then, $\sigma'(p)$ is calculated by evaluating the transverse equation at q .

The diffeomorphism invariance condition may be solved using a functional Newton’s method as follows. Manifolds near an initial manifold are represented by a certain Banach space B of embeddings of a fixed standard manifold, such as the standard 2-torus. The diffeomorphism $\tilde{F}$ induces an operator T on B whose fixed point corresponds to $\tilde{V}$. To solve for the fixed point of T , Newton’s method is used with the embedding of the initial manifold as initial data. A main issue here is to choose B so that T has a unique fixed point in B .

The vector field invariance condition may be solved using a functional Newton’s method as follows. Manifolds near an initial manifold are represented by embeddings in $\mathbb{R}^n$ of a fixed standard manifold S . The vector field invariance condition is formulated as a first order PDE, where the independent variables are coordinates on S , and the dependent variable $u \in \mathbb{R}^n$ gives a parametrization of the perturbed manifold. A parametrization of a manifold is not unique. Thus, extra constraints on the parametrization u are added. The resulting PDE is discretized and solved using Newton’s method. A main issue here is the stability and convergence of the nonlinear PDE discretization.

A computer algorithm requires a formulation where all the components have a finite representation. In the above approaches, smooth manifolds are replaced by *discrete manifolds*. These are C^0 , typically Lipschitz, manifolds which approximate smooth manifolds. These discrete manifolds are computable from data consisting of a set of discrete points called *grid points*. The grid points are typically points in the discrete manifold, with some additional structure. During a run of an algorithm, a sequence of discrete manifolds converges to a discrete manifold $\tilde{V}_D$ which approximates the exact perturbed manifold $\tilde{V}$. One *iterate* of the algorithm determines the next successive approximation to $\tilde{V}_D$. When applicable, a functional Newton method provides quadratic convergence of this iteration. The graph transform converges linearly.

Figure 1. Lorenz system orbit and hyperbolic splitting; two tori in the Lorenz-84 system, moving away from a Hopf saddle-node bifurcation [47].

In general, a numerical implementation of a functional Newton's method requires the solution of a large linear system at each iterate. The size of the system depends on the size of the discretizing grid. This is because the functional Newton's method solves for all the grid points in the discrete manifold simultaneously. For quasiperiodic tori and the diffeomorphism invariance condition, this numerical issue has been treated in [23]. This is done using either the normal hyperbolicity or by imposing extra non-resonance conditions. On the other hand, the graph transform solves for each grid point on the manifold independently of other grid points. This also makes the graph transform ideally suited to a parallel computer implementation, as remarked in [12].

1.2.2. Specific results. Next, some of the existing literature on the problem of computing a perturbed manifold is indicated. Usually, one of the three general approaches described in section 1.2.1 is used.

For quasiperiodic tori, the diffeomorphism condition is solved using a functional Newton's method in [22–24]. Here, the Banach space B consists of global parametrizations of tori where the parametrizations conform to the restricted dynamics. This restricted dynamics is a given rotation. Convergence of Newton's method is proved under the hypothesis of normal hyperbolicity. There is empirical evidence of convergence in the normally elliptic case. Precedents include [8, 48]. The computation of quasiperiodic tori is addressed by solving the vector field invariance condition using a functional Newton's method in [45, 46].

For attracting tori of vector fields, an algorithm based on the graph transform is given in [12]. There is a proof of convergence under the hypothesis of normal hyperbolicity. The underlying equations of the graph transform are solved using a shooting method to solve the associated boundary value problems. See also [41] for an extension of this work. The computation of attracting tori is addressed by solving the vector field invariance condition using a functional Newton's method in [11]. This is done under a hypothesis stronger than normal hyperbolicity. See also [36] for an extension of this work.

For attracting or saddle type manifolds with trivial normal bundle, an algorithm based on the graph transform is given in [4, 37]. The underlying equations of the graph transform are solved by replacing the perturbed restricted equation with the unperturbed version. Then, for example, in the attracting case, the two underlying equations of the graph transform (the restricted and transverse equations) are combined and solved together using Newton's method. For saddle type manifolds, the invariant hyperbolic splitting is computed using the linear graph transform.

An alternative approach to computing invariant manifolds is presented in [18, 19]. The vector field invariance condition is solved in a point-wise manner, in contrast to the functional

Newton's method given above. The focus is on attracting manifolds of dissipative systems in chemical kinetics and statistical physics. See also [17].

1.3. Discretizations of manifolds

Given a method of solving for an invariant manifold such as those discussed above, there is often considerable leeway in how to discretize the manifold. Methods of discretizing manifolds include global parametrizations of tori (for example, by truncated generalized Fourier series [22]), splines [41], finite element methods, simplicial complexes [4] or a discrete tubular neighbourhood as in the present paper. Simplicial complexes and other structured grids of points approximating a manifold are the simplest and most generally applicable methods of approximating manifolds. Spline and finite element approximations give higher order approximations and/or greater global smoothness and are also in principle generally applicable. A discrete tubular neighbourhood allows a manifold to be approximated by a graph, see section 1.4.

For several runs of the algorithm, the discrete approximations $\tilde{V}_D^1, \tilde{V}_D^2, \dots$ converge to the exact perturbed manifold $\tilde{V}$ as the mesh size of the discrete manifold $\rightarrow 0$. For example, the order of convergence is 1 for simplicial complexes and of higher order ≥ 1 for splines, finite element methods or the discrete tubular neighbourhood approach of the present paper. The convergence of a single run of an algorithm to $\tilde{V}_D$ can also be effected by the type of discretization of the manifold. Section 4.2 contains an illustration of this for the graph transform.

1.4. The graph transform and the present approach

The algorithm of the present paper is adapted from one of the classical approaches to the proof of the invariant manifold theorem, the graph transform. The theory of invariant manifolds using the graph transform is well developed [27]. As discussed in section 1.2.1, the graph transform operates on sections $\sigma : V \rightarrow N(V)$, where the vector bundle $N(V)$ is associated with a tubular neighbourhood of the submanifold V . To implement the graph transform, discrete sections and a discrete tubular neighbourhood of V are constructed. Since a tubular neighbourhood of V is the geometrical setting of the graph transform, a discrete tubular neighbourhood is a natural approach which allows an analogous development of a discrete graph transform. In particular, the convergence properties of the graph transform are inherited by the algorithm. This complete theory of convergence is one thing that distinguishes the present approach from many other approaches to computing invariant manifolds in the literature.

A discrete tubular neighbourhood is induced by a (Lipschitz) transverse vector bundle N . The base space of the discrete vector bundle N is the polyhedron P of a simplicial complex supporting V . The vector bundle N is constructed by interpolating k -plane fibres. This is done by combining simplex-wise affine interpolation with an explicit covering of the Grassman manifold. Manifolds near P are approximated by graphs of discrete sections of N . The discrete sections are locally polynomial maps from a face of P to the transverse bundle. A globally continuous approximation by local polynomial graphs of arbitrary order is obtained for manifolds near P .

1.4.1. Graph transform solution strategy. In section 5.2, a practical approach to solving the global equations associated with the graph transform is proposed. For example, in the attracting case, the restricted equation must be solved for a point $p \in P$, for each grid point. The most

potentially work intensive step of the running algorithm is determining a neighbourhood U in P containing p . It is necessary to find U only for the first iterate, since for later iterates data from previous iterates can be used.

Once U is known, obtaining p to a desired tolerance is quickly accomplished using a standard root finding method. In fact, the amount of numerical work involved in an iterate of the graph transform after the first iterate is equal to solving a d -dimensional equation (the restricted equation) for a local root using a standard root finding method, with a good initial guess, m times. Here, m is the number of grid points and $d = \dim V$. (There may also be numerical work associated with modifications to the transverse bundle and redistribution of the grid points—this is done between runs of the graph transform algorithm. See section 5.3.)

In the algorithm proposed in section 5.2, for each grid point, U is determined using a simple geometric test on each d -simplex of P . This is an efficient method of determining an initial guess for the local root finding method. Thus, although the convergence of the graph transform is linear, this makes the computation of each iterate of the graph transform fast.

1.4.2. Continuation and the linear graph transform. So far, the focus has been on the problem of computing an $\tilde{F}$ -invariant manifold $\tilde{V}$, given an F -invariant manifold V , for a perturbed system $\tilde{F}$ near F . To compute a manifold of a system which is not near a system with a known initial manifold V , the algorithm may be repeated with computed initial data. Thus, the problem of computing a perturbed manifold, discussed so far, is the main repeated step in a continuation process.

In the general saddle case, the initial data for the graph transform algorithm consist of an F -invariant manifold V together with a splitting of $N(V)$ into stable and unstable bundles, $N(V) = N^u(V) \oplus N^s(V)$. Thus, one step of the continuation process has two parts. The initial data are an F -invariant manifold V with a splitting $N(V) = N^u(V) \oplus N^s(V)$. The first step uses the graph transform Γ , constructed in terms of V and its hyperbolic splitting, to determine the $\tilde{F}$ -invariant manifold $\tilde{V}$. The second step uses linear graph transforms $\mathcal{L}^s$ and $\mathcal{L}^u$ together with initial data determined by $\tilde{V}$ and the hyperbolic splitting of V to determine the hyperbolic splitting of $\tilde{V}$. Here, $\mathcal{L}^s$ determines $N^s(\tilde{V})$ and $\mathcal{L}^u$ determines $N^u(\tilde{V})$. Now the first and second steps are repeated with initial data given by $\tilde{V}$ and its hyperbolic splitting.

For example, $\mathcal{L}^u$ essentially operates on a certain space of vector bundles, using the linearization of $\tilde{F}$ to obtain successive approximations to $N^u(\tilde{V})$. To discretize $\mathcal{L}^u$, discrete vector bundles are constructed by interpolating k -plane fibres. The discrete vector bundle allows a derivation of the discrete linear graph transform analogous to the usual derivation of the linear graph transform. Thus, as for the graph transform, the convergence properties of the linear graph transforms are inherited by their discretizations.

1.4.3. Summary. Compared with related work, the present approach gives a general purpose algorithm. It applies to manifolds of arbitrary topological type, attracting or saddle type, regardless of the restricted dynamics. There is a satisfactory theory of convergence in this general setting. In addition, the numerical approach for solving the underlying global equations of the graph transform is simple and motivated by the geometry of the graph transform. If the manifold is not normally hyperbolic, however, a different approach should be used, for example, [22].

1.5. Outline of paper

In section 2.1, normal hyperbolicity, the invariant manifold theorem, tubular neighbourhoods and hyperbolic splittings are discussed. In section 2.2, the domain of Γ , a certain space of Lipschitz sections $V \rightarrow N(V)$, is specified and Γ is formulated starting from an invariance condition. In section 2.3, the linear graph transforms $\mathcal{L}^s$ and $\mathcal{L}^u$ along with their domains are formulated. In these sections, the contraction mapping properties of Γ , $\mathcal{L}^s$ and $\mathcal{L}^u$ are discussed. Section 2 is close in spirit to [27].

In section 3, the discretizations of the domains of Γ , $\mathcal{L}^s$ and $\mathcal{L}^u$ are formulated. To do this, a discrete tubular neighbourhood along with a space of discrete sections of the associated vector bundle is constructed. In section 3.1, the polyhedral approximation P to V is described. In sections 3.2 and 3.3, discrete vector bundles $N^u(P)$ and $N^s(P)$ approximating $N^u(V)$ and $N^s(V)$ are constructed. In section 3.4, discrete approximating sections of $N^u(P) \oplus N^s(P)$ are constructed.

In section 4, discrete versions of Γ , $\mathcal{L}^s$ and $\mathcal{L}^u$ are formulated, based on the constructions of section 3. In sections 4.1 and 4.2, the discrete sections described in section 3.4 lead to a discrete version Γ_D of the graph transform. Starting from the formulation of section 2.2, Γ_D is constructed in two steps. In section 4.1, the smooth transverse bundle $N(V)$ is replaced by the Lipschitz transverse bundle $N(P)$. The graph transform defined in this way is a contraction on a certain space of sections $\mathcal{S}$ whose fixed point gives the exact perturbed manifold $\tilde{V}$. In section 4.2, the space $\mathcal{S}$ of section 4.1 is replaced by the subset of $\mathcal{S}$ consisting of discrete sections. The discrete graph transform Γ_D defined in this way is a contraction whose fixed point approximates $\tilde{V}$. In section 4.3, the discrete approximations of the stable and unstable bundles given in section 3.3 lead to discrete versions $\mathcal{L}_D^s$ and $\mathcal{L}_D^u$ of the linear graph transforms. The discrete linear graph transforms are contractions on certain spaces of sections whose fixed points approximate $N^s(\tilde{V})$ and $N^u(\tilde{V})$.

In section 5, an outline of a computer implementation of the formulated algorithm is given. In section 5.1, the general continuation context is discussed. In section 5.2, a computer implementation of the graph transform is given. The main inner loop is described in section 5.3. In section 5.3, numerical issues that appear to be generally relevant are addressed. These include numerical conditioning of the underlying equations of the graph transform as well as smoothing techniques for both the discrete vector bundles $N^u(P)$ and $N^s(P)$ and discrete sections of the transverse bundle $N^u(P) \oplus N^s(P)$.

Section 6.1 contains some of the numerical experiments performed to validate the algorithm. These show that the empirical performance of the algorithm matches the theoretical expectations. Section 6.2 contains an application to computing the ‘slow transient’ surface of an enzyme reaction model. This application requires a modification to the main algorithm applied in section 6.1. This is because what is computed in this case is just part of the invariant surface in the region of interest and not the whole invariant surface. See also the DISC project website <http://home.nethere.net/hagen>.

Remarks

- (i) An important practical feature of the present approach is that it works regardless of the dynamics on the invariant manifold. For example, $\tilde{V}$ may contain a smaller attractor. This is a case in which simple iteration of the dynamical system, sometimes used, fails to adequately approximate $\tilde{V}$. Simple iteration fails for another reason if $\tilde{V}$ is of the saddle type. Here, numerical errors grow exponentially in the unstable direction. The present approach computes $\tilde{V}$ in a numerically stable way.

- (ii) The discrete tubular neighbourhood does in a computational setting what the tubular neighbourhood does for the graph transform. Namely, it uses V to split the problem of finding $\tilde{V}$ into two problems of lower dimension. One consequence of this is the following. Recall that P is the polyhedron of a simplicial complex supporting V . In appropriate coordinates, approximations to $\tilde{V}$ are locally *graphs* of polynomials $\{d\text{-simplex of } P\} \rightarrow \mathbb{R}^{n-d}$ rather than images of polynomials $\{d\text{-simplex of } P\} \rightarrow \mathbb{R}^n$, where $d = \dim V$ and $n = \dim M$. The discrete tubular neighbourhood allows us to piece together these polynomial graphs to approximate $\tilde{V}$ as a continuous global graph over V . This method of approximation is not restricted to the context of the graph transform. It may be suitable in other areas for computing successive approximations of manifolds which are not global graphs over a flat space.
- (iii) One main application of invariant manifolds is in bifurcation theory. Here, they appear both as a dimension-reducing mechanism in phase plus parameter space and as the manifestation of a bifurcation in phase space for a fixed non-critical parameter value. For example, the Hopf bifurcation yields a periodic orbit on a surface in phase plus parameter space. In this context, the current algorithm does not address the possibility that, while varying the continuation parameter, a bifurcation may occur which manifests in the coexistence of two or more nearby normally hyperbolic manifolds. In this case, there is an intermediate parameter value at which the normal hyperbolicity is lost. In the present paper it is assumed that $\tilde{V}$ is the unique invariant manifold Lipschitz-near V , as in the setting of the invariant manifold theorem.

2. Invariant manifolds

In this section, normally hyperbolic invariant manifolds are introduced. An overview of some basic definitions and results from [27] is given. The graph transform is set up as a way of locating both the perturbed manifold and the perturbed hyperbolic splitting. Later, the elements of these constructions, for example, tubular neighbourhoods and sections of vector bundles, will be specified in a numerical implementation.

2.1. Normal hyperbolicity

Our starting point is a C^r diffeomorphism F on a C^∞ Riemannian manifold M , with an invariant submanifold $V \subset M$. Here, V is a compact, C^r , r -normally hyperbolic submanifold of M , $r \geq 1$. Recall that V is r -normally hyperbolic for F if there is a DF -invariant splitting

$$T_V(M) = N^u(V) \oplus T(V) \oplus N^s(V) \quad (1)$$

and a Riemann structure on the tangent bundle $T_V(M)$, such that, for $y \in V$, $i \geq 0$ and $0 \leq k \leq r$:

$$\begin{aligned} \|DF^i|_{N_y^s(V)}\| \cdot \|(DF^i|_{T_y(V)})^{-1}\|^k &\leq c\mu^i, \\ \|(DF^i|_{N_y^u(V)})^{-1}\| \cdot \|DF^i|_{T_y(V)}\|^k &\leq c(1/\lambda)^i, \end{aligned} \quad (2)$$

for some $0 < \mu < 1 < \lambda < \infty$ and $0 < c < \infty$. Here the operator norms are associated with the Riemann structure on $T_V(M)$. Translated into words, condition (2) is the following: V is r -normally hyperbolic for F if, along trajectories of F on V , DF^i expands $N^u(V)$ more than DF^{ki} expands $T(V)$ and DF^i contracts $N^s(V)$ more than DF^{ki} contracts $T(V)$, $0 \leq k \leq r$, for all i sufficiently large.

According to the *invariant manifold theorem* [27, theorem 4.1], a C^r diffeomorphism $\tilde{F}$, that is C^r -near F , has an r -normally hyperbolic invariant manifold $\tilde{V}$, that is C^r and C^r -near V . This suggests the possibility of computing an invariant manifold by numerical continuation. That is, given the homotopic diffeomorphisms $\tilde{F}$ and F as well as V successively increment the homotopy parameter, computing an approximation to the perturbed manifold at each step, to obtain $\tilde{V}$. To implement this idea, we look at the proof of the invariant manifold theorem.

To linearize $\tilde{F}$ at V , a tubular neighbourhood of V is used [26, 31]. A tubular neighbourhood of V in M is a vector bundle E with base space V , an open neighbourhood U of V in M , an open neighbourhood Z of the zero section in E and a homeomorphism $\phi : Z \rightarrow U$. Here, ϕ must satisfy $\phi \circ \sigma_0 = i$, where $\sigma_0 : V \rightarrow E$ is the zero section and $i : V \rightarrow M$ is the inclusion. For example, the normal bundle $E = \bigcup_{p \in V} T_p(V)^\perp$ of V in M gives a tubular neighbourhood of V , at least if $r \geq 2$. In fact, any Lipschitz vector bundle $N(V)$, transverse to $T(V)$ in $T_V(M)$, gives a tubular neighbourhood of V in M . In what follows, $\tilde{V}$ is constructed in the neighbourhood U in M or equivalently in the neighbourhood Z in $N(V)$. Here, it is appropriate for Z to be the closure of a neighbourhood, $Z = Z(\epsilon) = \{(p, v) \in N(V) : |v|_p \leq \epsilon\}$.

The graph transform uses the $\tilde{F}$ -dynamics near V to locate $\tilde{V}$. In the attracting case, when $N_y^u(V) = \{0\}$, $y \in V$, iterating the graph transform boils down to taking successive $\tilde{F}$ images of V . In the saddle case, when both $N^s(V)$ and $N^u(V)$ are non-trivial, the problem is split into two parts. For any Lipschitz transverse vector bundle $N(V)$, the invariant splitting (1) induces a splitting of $N(V) = N^u(V) \oplus N^s(V)$ into stable and unstable parts. The splitting $T_V(M) = N^u(V) \oplus T(V) \oplus N^s(V)$ has the same growth properties (2) as the invariant splitting. Sections of Z may now be written as $\sigma(p) = (p, v^s(p), v^u(p))$, where $v^s(p) \in Z_p^s = N_p^s(V) \cap Z$, $v^u(p) \in Z_p^u = N_p^u(V) \cap Z$.

2.2. The graph transform

Effectively, the graph transform takes as input a manifold near V and returns as output a manifold near V . Graphs of sections of the vector bundle $Z = Z(\epsilon)$, discussed above, are these manifolds near V . Thus, the domain of the graph transform is a certain space of sections of the vector bundle $Z = Z(\epsilon)$. In fact, the graph transform is a contraction on a space of Lipschitz sections $\sigma : V \rightarrow Z$. To define the Lipschitz constant of a section, a C^0 connection in $T_V(M)$ is used [34]. A connection gives a way of comparing points in different fibres of $T_V(M)$. It does this using a continuous family of horizontal subspaces $H(y)$, $y \in T_V(M)$, which extend the tangent spaces of V . More precisely, a C^0 connection in the vector bundle $\pi : T_V(M) \rightarrow V$ is a C^0 distribution $H : T_V(M) \rightarrow T(T_V(M))$ with $T_y(T_V(M)) = H(y) \oplus V(y)$, $y \in T_V(M)$, where $V(y)$ is the kernel of $D\pi$. Here, it is also required that the horizontal subspace of the associated frame bundle corresponding to $H(y)$ be invariant under the structure group. This implies, in particular, that if $\sigma_0 : V \rightarrow T_V(M)$ is the zero section, then $H(\sigma_0(p)) = D\sigma_0(T_p(V))$.

To define the slope of a section $\sigma : V \rightarrow T_V(M)$ at $p \in V$, let $\theta : V \rightarrow T_V(M)$ be a C^1 section with $\theta(p) = \sigma(p)$ and $D\theta(T_p(V)) = H(\sigma(p))$. Then the slope of σ at p is

$$\text{slope}_p(\sigma) = \limsup_{x \rightarrow p} \frac{|\sigma(x) - \theta(x)|_x}{d_V(x, p)}$$

[27]. Since Z^s and Z^u are subbundles of $T_V(M)$, this also gives a natural definition of the slope of sections $\sigma^s : V \rightarrow Z^s$ and $\sigma^u : V \rightarrow Z^u$. From this, the Lipschitz constant of σ^s is $\text{Lip}(\sigma^s) = \sup_{p \in V} \text{slope}_p(\sigma^s)$, and similarly for σ^u . Now, the Lipschitz constant of a section $\sigma(p) = (p, v^s(p), v^u(p))$ of Z is $\text{Lip}(\sigma) = \max\{\text{Lip}(\sigma^s), \text{Lip}(\sigma^u)\}$, where

Figure 2. Invariance condition (3).

$\sigma^s(p) = (p, v^s(p))$ and $\sigma^u(p) = (p, v^u(p))$. The domain of the graph transform is $\mathcal{S}_{\epsilon, \delta} = \{\sigma : V \rightarrow Z : \text{Lip}(\sigma) \leq \delta\}$. The norm on $\mathcal{S}_{\epsilon, \delta}$ is $\|\sigma\| = \max\{|\sigma^s|_s, |\sigma^u|_u\}$, where $|\cdot|_s$ and $|\cdot|_u$ are the natural C^0 norms on sections of Z^s and Z^u , respectively. With this norm, $\mathcal{S}_{\epsilon, \delta}$ is complete. In fact, the graphs of the sections in $\mathcal{S}_{\epsilon, \delta}$ are the Lipschitz manifolds near V in the Lipschitz norm.

To formulate the graph transform, the starting point is the $\tilde{F}$ -invariance condition $\phi \circ \sigma(V) = \tilde{F} \circ \phi \circ \sigma(V)$. This is split into two coupled equations, a part on V and a part normal to V . We put $\tilde{F}^0 = \phi^{-1} \circ \tilde{F} \circ \phi$ and work in $N(V)$. The image of $\phi \circ \sigma$ is $\tilde{F}$ -invariant if and only if

$$\begin{aligned} (y, v^s(y), v^u(y)) &= \tilde{F}^0(p, v^s(p), v^u(p)), \\ y &= \pi \circ \tilde{F}^0(p, v^s(p), v^u(p)), \end{aligned} \quad (3)$$

for $p \in V$, where $\pi : N(V) \rightarrow V$ is the vector bundle projection. See figure 2. Under our hypotheses, $y = \pi \circ \tilde{F}^0(p, v^s(p), v^u(p))$ may be solved for a unique $p \in V$ given $y \in V$ and $\sigma \in \mathcal{S}_{\epsilon, \delta}$ for small ϵ, δ and $\theta = \|F - \tilde{F}\|_{C^1}$. Denote this solution by $p = p(y, v^s, v^u)$. Now, given $\sigma \in \mathcal{S}_{\epsilon, \delta}$, $\sigma(p) = (p, v^s(p), v^u(p))$, the *graph transform* of σ is the section $\Gamma(\sigma)(p) = (p, w^s(p), w^u(p))$. Here, w^s is defined by

$$w^s(y) = P_y^s \circ \tilde{F}^0(p, v^s(p), v^u(p)), \quad p = p(y, v^s, v^u), \quad (4)$$

for $y \in V$, where $P_y^s : N_y(V) \rightarrow N_y(V)$ is the linear projection with range $N_y^s(V)$ and nullspace $N_y^u(V)$. The unstable part w^u is defined implicitly by

$$\begin{aligned} v^u(y) &= P_y^u \circ \tilde{F}^0(p, v^s(p), w^u(p)), \\ y &= \pi \circ \tilde{F}^0(p, v^s(p), w^u(p)), \end{aligned} \quad (5)$$

for $p \in V$, where $P_y^u : N_y(V) \rightarrow N_y(V)$ is the linear projection with range $N_y^u(V)$ and nullspace $N_y^s(V)$. In (5), there is a unique solution for $w^u(p)$ for small θ, ϵ and δ .

If $\sigma = \Gamma(\sigma)$, then (4) and (5) imply (3). Hence σ is a fixed point of Γ if and only if the graph of σ is $\tilde{F}$ -invariant. By replacing $\tilde{F}$ with $\tilde{F}^N$ above, for some large integer N , Γ becomes a contraction on $\mathcal{S}_{\epsilon, \delta}$ whose fixed point σ^* satisfies $\phi \circ \sigma^*(V) = \tilde{V}$.

2.3. The linear graph transform

One step of the proposed continuation algorithm has two parts. The initial data are an F -invariant manifold V with splitting $N^u(V) \oplus T(V) \oplus N^s(V)$. The first step uses the graph transform Γ on V with $N^u(V) \oplus T(V) \oplus N^s(V)$ to determine the $\tilde{F}$ -invariant manifold $\tilde{V}$. That is, starting with the zero section σ_0 , Γ is iterated, $\Gamma^i(\sigma_0) \rightarrow \sigma^*$ in C^0 norm as $i \rightarrow \infty$. The second step uses linear graph transforms $\mathcal{L}^s$ and $\mathcal{L}^u$ together with initial data determined by $\tilde{V}$ and $N^u(V) \oplus T(V) \oplus N^s(V)$ to determine the splitting $N^u(\tilde{V}) \oplus T(\tilde{V}) \oplus N^s(\tilde{V})$ for $\tilde{V}$. Then the first and second steps are repeated with initial data $\tilde{V}$, $N^u(\tilde{V}) \oplus T(\tilde{V}) \oplus N^s(\tilde{V})$.

To illustrate the details of the linear graph transform, here $\mathcal{L}^u$ is formulated. Given a transverse bundle $N(\tilde{V})$, first the initial data for $\mathcal{L}^u$ in $N(\tilde{V})$ is determined. Let

$Q : T_{\tilde{V}}(M) \rightarrow T_{\tilde{V}}(M)$ be, on each fibre $T_y(M)$, the linear projection with range $N_y(\tilde{V})$ and nullspace $T_y(\tilde{V})$. Initial data $N(\tilde{V}) = N^{u,0}(\tilde{V}) \oplus N^{s,0}(\tilde{V})$ are then

$$N^{u,0}(\tilde{V}) = Q(N^{u,1}(\tilde{V})), \quad N^{s,0}(\tilde{V}) = Q(N^{s,1}(\tilde{V})),$$

where $N_y^{u,1}(\tilde{V})$, $N_y^{s,1}(\tilde{V})$ are obtained from $N_p^u(V)$, $N_p^s(V)$, $y = \phi \circ \sigma^*(p)$, by parallel translation $T_p(M) \rightarrow T_y(M)$ along the ϕ -images of fibres of $N(V)$, [1, 34]. There exists $\alpha > 0$, where $\alpha \rightarrow 0$ as $\epsilon + \delta + \theta \rightarrow 0$, such that, if $\{\angle N(V), T(V)\}$, $\{\angle N(\tilde{V}), T(\tilde{V})\} \geq \alpha > 0$, then this procedure produces non-degenerate initial data $N^{u,0}(\tilde{V})$, $N^{s,0}(\tilde{V})$.

Now, the domain of $\mathcal{L}^u$ is a space of sections whose graphs are the j -plane bundles near $N^{u,0}(\tilde{V})$ in $N(\tilde{V})$, where j is the dimension of $N^{u,0}(\tilde{V})$. These are sections of the bundle $L(\tilde{V})$ whose fibre at $y \in \tilde{V}$ is the space of linear transformations $N_y^{u,0}(\tilde{V}) \rightarrow N_y^{s,0}(\tilde{V})$, $L(N_y^{u,0}(\tilde{V}), N_y^{s,0}(\tilde{V}))$, [27]. The domain of $\mathcal{L}^u$ is $\mathcal{S}_\eta = \{\sigma : \tilde{V} \rightarrow L(\tilde{V}) : \sup_y \|\sigma(y)\| \leq \eta\}$, where the operator norm $\|\cdot\|$ is associated with the Riemann structure on $T_{\tilde{V}}(M)$. The space $\mathcal{S}_\eta$ is complete with respect to the norm $|\sigma| = \sup_y \|\sigma(y)\|$.

To formulate $\mathcal{L}^u$, the starting point is the invariance condition. The linear mapping induced by $D\tilde{F} : T_{\tilde{V}}(M) \rightarrow T_{\tilde{V}}(M)$ on $N(\tilde{V}) \subset T_{\tilde{V}}(M)$ is $\Phi = Q \circ D\tilde{F}|_{N(\tilde{V})} : N(\tilde{V}) \rightarrow N(\tilde{V})$. The graph of $\sigma \in \mathcal{S}_\eta$ is Φ -invariant if and only if $\Phi(\text{graph}\{\sigma(x)\}) = \text{graph}\{\sigma(y)\}$, $y = \tilde{F}(x)$, $x \in \tilde{V}$. This condition is split into a part in $N^{u,0}(\tilde{V})$ and a part in $N^{s,0}(\tilde{V})$. Let $P_y^u : N_y(\tilde{V}) \rightarrow N_y(\tilde{V})$ be the linear projection with range $N_y^{u,0}(\tilde{V})$ and nullspace $N_y^{s,0}(\tilde{V})$. Define P_y^s analogously. Then the graph of $\sigma \in \mathcal{S}_\eta$ is Φ -invariant if and only if

$$\begin{aligned} \sigma(y)(\tilde{\rho}) &= P_y^s \circ \Phi(\rho, \sigma(x)(\rho)), \\ \tilde{\rho} &= P_y^u \circ \Phi(\rho, \sigma(x)(\rho)), \end{aligned} \tag{6}$$

for $\rho \in N_x^{u,0}(\tilde{V})$, $x \in \tilde{V}$, where $y = \tilde{F}(x)$. The second equation in (6) is a linear mapping $N_x^{u,0}(\tilde{V}) \rightarrow N_y^{u,0}(\tilde{V})$, $\rho \rightarrow \tilde{\rho}$, which is invertible for small ϵ , δ , θ and η . Denote the inverse $B_y(\tilde{\rho}) = \rho$. Then, the graph transform of σ is the section $\mathcal{L}^u(\sigma)(y) = P_y^s \circ \Phi \circ (\text{id}, \sigma(x)) \circ B_y$ for $y \in \tilde{V}$. Here, $(\text{id}, \sigma(x)) : N_x^{u,0}(\tilde{V}) \rightarrow N_x(\tilde{V})$ is $(\text{id}, \sigma(x))(\rho) = (\rho, \sigma(x)(\rho))$.

The graph of σ is Φ -invariant if and only if σ is a fixed point of $\mathcal{L}^u$. By replacing Φ with Φ^N above, for some large integer N , and for ϵ , δ , θ and η small, $\mathcal{L}^u$ is a contraction on $\mathcal{S}_\eta$ whose fixed point σ^* gives the Φ -invariant bundle $N^u(\tilde{V})$. More precisely, the graph of σ^* is the desired unstable part $N^u(\tilde{V})$ of the splitting of $N(\tilde{V})$. The formulation of $\mathcal{L}^s$ is analogous. Similar approaches work to determine, for example, the invariant transverse bundle.

3. Discrete sections

In this section and the next, a discrete version of the graph transform, realizable in a numerical implementation, is modelled on the formulation of section 2. The main part is the discretization of the candidate manifolds, the graphs of sections in $\mathcal{S}_{\epsilon,\delta}$. From this, a discrete version of the graph transform follows naturally. The discretization of the candidate manifolds is the subject of the present section. In this section and the next, the manifold $M = \mathbb{R}^n$ with the constant Riemann metric induced by the usual inner product. This is not, in principle, a reduction of

the generality of the method, since V may be embedded in $\mathbb{R}^n$ and the property of normal hyperbolicity (2) is independent of the Riemann structure.

As discussed in section 2, the proposed continuation algorithm requires certain initial data to start. The initial data consist of a C^r F -invariant submanifold $V \subset \mathbb{R}^n$, $r \geq 1$, with splitting $T_V(\mathbb{R}^n) = N^u(V) \oplus T(V) \oplus N^s(V)$, where the transverse bundle $N(V) = N^u(V) \oplus N^s(V)$ is Lipschitz. The goal is to compute, for the final value of the continuation parameter, the same data we started with: $\tilde{V}$ and its hyperbolic splitting. This enables the algorithm to be restarted. For a computer algorithm, discrete approximations must be made to the components of the graph transform. An approximation to V is given in section 3.1. An approximation to the splitting $N^u(V) \oplus T(V) \oplus N^s(V)$ is provided in sections 3.2 and 3.3. An approximation to a section in $\mathcal{S}_{\epsilon,\delta}$ is given in section 3.4.

3.1. Initialization

A finite approximation to V is provided by a geometric simplicial complex $\mathcal{C} \subset \mathbb{R}^n$ supporting $V \subset \mathbb{R}^n$, [6, 35, 51]. Recall that the polyhedron $P \subset \mathbb{R}^n$ of $\mathcal{C}$ is the set of all points in the simplices of $\mathcal{C}$ with the subspace topology. A simplicial complex $\mathcal{C}$ supports V if the vertices of all simplices are in V and P is homeomorphic to V . If H is the maximal diameter of the simplices of $\mathcal{C}$ then P converges to V in the Lipschitz norm as $H \rightarrow 0$. Denote by $\mathcal{C}_1 \dots \mathcal{C}_N$ the d -simplices of $\mathcal{C}$, $d = \dim V$. There is a restriction on the types of supporting simplicial complexes $\mathcal{C}$ that will be used. Namely, approximations of candidate manifolds will use polynomial maps of arbitrary order on each $\mathcal{C}_i$. For the uniformity of these approximations as $H \rightarrow 0$, it is required that $\{\mathcal{C}_i\}_{i=1}^N$ be a regular family. This means that if h_i is the diameter of $\mathcal{C}_i$ and ρ_i the supremum of the diameters of the inscribed spheres of $\mathcal{C}_i$, then h_i/ρ_i is bounded uniformly for all i and $H \rightarrow 0$, [9].

3.2. Stable, unstable and transverse bundles

An approximation of a candidate manifold is now obtained. This approximation is the graph of a discrete section of a transverse vector bundle associated with a tubular neighbourhood of P . Thus, a discrete tubular neighbourhood of P must be constructed first. The tubular neighbourhood enables us to identify candidate manifolds and their approximations, with graphs of sections over P . In the process of constructing the tubular neighbourhood, discrete approximations to the splitting $N^u(V) \oplus T(V) \oplus N^s(V)$ will also be constructed. In fact, the approximation to this splitting will be given by vector bundles $N^s(P)$ and $N^u(P)$, where $N(P) = N^s(P) \oplus N^u(P)$ is the transverse bundle associated with the tubular neighbourhood of P .

A tubular neighbourhood of P is induced by a transverse field of k_0 -planes $\mu : P \rightarrow \mathcal{G}_{n,k_0}$ = the Grassmann manifold of k_0 -planes of $\mathbb{R}^n$, $k_0 = \text{codim } V$, provided μ is locally Lipschitz with respect to Riemannian metrics [26, 52]. The approximation to the hyperbolic splitting satisfies $N_x(P) = N_x^s(P) \oplus N_x^u(P) \subset T_x(\mathbb{R}^n)$, $x \in P$. Here, $T_x(\mathbb{R}^n)$, $x \in P$, are as usual identified with the ambient space $\mathbb{R}^n$ containing V and also the underlying space $\mathbb{R}^n$ of the Grassmann manifold via the standard basis. By this identification, the field μ gives a transverse bundle $N(P)$. In fact, the field μ is made up of two parts, $\mu(x) = \mu_1(x) \oplus \mu_2(x)$, $x \in P$, where $\mu_i : P \rightarrow \mathcal{G}_{n,k_i}$ for $i = 1, 2$, $k_1 = \dim N^s(V)$ and $k_2 = \dim N^u(V)$. Here, μ_1 gives $N^s(P)$ and μ_2 gives $N^u(P)$.

Next, $N^s(V)$ is approximated by $N^s(P)$. To be precise, the given $N(V)$ induces a homeomorphism $\psi : P \rightarrow V$. Let $N^{s,0}(P)$ be the vector bundle over P whose fibre at $y \in P$

is $N_{\psi(y)}^s(V)$. To approximate $N^s(V)$, the Lipschitz field $\vartheta : P \rightarrow \mathcal{G}_{n,k_1}$, $\vartheta(y) = N_y^{s,0}(P)$, is approximated by a field $\mu_1 : P \rightarrow \mathcal{G}_{n,k_1}$. The field μ_1 is constructed by interpolating a given finite set of data points in $\mathcal{G}_{n,k_1}$. These data points are the k_1 -planes $\{N_y^s(V) : y \in \mathcal{C}^0\}$, where $\mathcal{C}^0$ is the set of vertices of $\mathcal{C}$.

Generally, interpolation is the process of approximating a function $f : A \rightarrow B$ using only a finite set of discrete data. Typically, f must have certain regularity (smoothness) properties. The discrete data set determines a constructible interpolant $\mathcal{I} \circ f : A \rightarrow B$ which approximates $f : A \rightarrow B$. For example, the discrete data set may consist of function values $b_i \in B$ at a discrete set of points $a_i \in A$, $f(a_i) = b_i$. An interpolant may be fitted to this data, $\mathcal{I} \circ f(a_i) = b_i$. In the present case, the Lipschitz field $\vartheta : P \rightarrow \mathcal{G}_{n,k_1}$, $\vartheta(y) = N_y^{s,0}(P)$, is approximated by an interpolant $\mu_1 : P \rightarrow \mathcal{G}_{n,k_1}$.

In the case $k = 1$ the following method may be used to interpolate k -fibres over a simplex. Given $d + 1$ nearby 1-plane fibres at the vertices of a d -simplex, choose $d + 1$ unit vector bases $b_1 \dots b_{d+1}$ for the fibres, all contained in a small neighbourhood in the frame manifold. Then a basis for the interpolating 1-plane fibre at the barycentric coordinates $(t_1, \dots, t_{d+1})$ is obtained by normalizing the vector $v = t_1 \cdot b_1 + \dots + t_{d+1} \cdot b_{d+1}$. This is numerically practical since the nearness of the bases $b_1 \dots b_{d+1}$ implies that $|v|$ is near one. This approach may be extended to the case $k > 1$ as well [4]. In general, there are two parts to the problem of interpolating k -fibres over a simplex. The first is choosing bases for the fibres and the second is interpolating the fibre bases. For example, given $d + 1$ nearby k -plane fibres at the vertices of a d -simplex and an ordered orthonormal basis $\mathcal{B}$ for one fibre, obtain bases for the other k -plane fibres by orthogonal projection of $\mathcal{B}$ onto those fibres. Then to obtain a basis for the interpolating k -plane fibre at the barycentric coordinates $(t_1, \dots, t_{d+1})$, use the $k = 1$ procedure for each basis vector. To obtain an orthonormal basis for the interpolating fibre, an orthogonalization procedure is applied to this basis.

Section 3.3 below is motivated by a desire to find a natural way to interpolate k -fibres over a simplex, $k \geq 1$. For one thing, it is not necessary to leave the realm of orthonormal frames to do this. For example, in the $k = 1$ case given above, it is possible to interpolate the bases in such a manner that the basis vector v remains a unit vector, while the interpolation parameters $(t_1, \dots, t_{d+1})$ vary. In the approach described below, plane rotation matrices are used to interpolate special orthonormal bases for the k -plane fibres. Interpolating the fibres involves multiplication of the orthonormal bases by orthogonal matrices. In contrast to the $k = 1$ approach of the previous paragraph, no normalization of non-unit vectors is required. There is no point in the procedure at which the bases are not orthonormal. This is numerically desirable. In addition, this approach in the case $k > 1$ is hardly more algorithmically complicated than the $k = 1$ case.

3.3. Constructing bundles from data points

The field μ_1 is constructed from data points in $\mathcal{G}_{n,k_1}$ at the vertices of P . To define μ_1 on P , these data points are interpolated over each d -simplex $\mathcal{C}_i$ in a consistent manner. The interpolation is performed in a certain space of orthonormal frames for the k_1 -planes of $\mathcal{G}_{n,k_1}$. Since the same procedure is used for μ_2 , in the following we will use k to denote a variable which may be k_1 or k_2 . Recall that $\mathcal{F}_{n,k}$, the space of k -frames in $\mathbb{R}^n$, $k \leq n$, is given the structure of a smooth manifold by its natural identification with the space of $n \times k$ matrices of rank k . The space of $n \times k$ matrices of rank k is a smooth manifold due to its identification with an open subset of $\mathbb{R}^{nk}$, [1].

The following lemma is used to choose an orthonormal frame for each k -plane data point in $\mathcal{G}_{n,k}$.

Lemma 1. Any k -plane $\mathcal{P} \in \mathcal{G}_{n,k}$ has an orthonormal frame $\rho_1 \dots \rho_k$ where the $n \times k$ matrix $[\rho_1 \dots \rho_k]$ is a lower triangular,

$$[\rho_1 \dots \rho_k] = \begin{bmatrix} * & 0 & \dots & 0 \\ * & * & \ddots & \vdots \\ \vdots & \vdots & \ddots & 0 \\ * & * & \dots & * \\ \hline & & & * \end{bmatrix}_{n \times k}, \quad (7)$$

and in addition $[\rho_1 \dots \rho_k]^T$ is in echelon form, for example

$$[\rho_1 \dots \rho_k]^T = \begin{bmatrix} * & * & * & * & * & * & * \\ 0 & * & * & * & * & * & * \\ 0 & 0 & 0 & * & * & * & * \\ 0 & 0 & 0 & 0 & 0 & 0 & * \end{bmatrix}_{4 \times 7}.$$

Proof. Choose an orthonormal frame $b_1 \dots b_k \in \mathbb{R}^n$ for $\mathcal{P}$. Orthogonally reduce $[b_1 \dots b_k]^T$ to upper triangular echelon form U [50], using, for example, a Givens reduction [16]. Thus, $O[b_1 \dots b_k]^T = U_{k \times n}$, where O is an orthogonal $k \times k$ matrix. Then $U^T = [b_1 \dots b_k]O^T$ is lower triangular and its columns form the desired orthonormal frame $\rho_1 \dots \rho_k$ for $\mathcal{P}$. $\square$

A lower triangular orthonormal k -frame in $\mathbb{R}^n$ is an orthonormal frame $\rho_1 \dots \rho_k$ where the $n \times k$ matrix $[\rho_1 \dots \rho_k]$ is lower triangular (7). The space of lower triangular orthonormal k -frames in $\mathbb{R}^n$ is a submanifold $\mathcal{T}_{n,k} \subset \mathcal{F}_{n,k}$. A reduced orthonormal frame is a lower triangular orthonormal frame $\rho_1 \dots \rho_k$ where $[\rho_1 \dots \rho_k]^T$ is in echelon form, as described in lemma 1. Given a k -plane $\mathcal{P}$, any two orthonormal frames $\rho_1 \dots \rho_k$ and $\rho'_1 \dots \rho'_k$ of $\mathcal{P}$ are always related by an orthogonal $k \times k$ matrix O , $[\rho'_1 \dots \rho'_k] = [\rho_1 \dots \rho_k]O$. If these are reduced orthonormal frames, then O is a diagonal matrix, whose diagonal entries are ± 1 . Hence, given a k -plane $\mathcal{P}$, a reduced orthonormal frame $\rho_1 \dots \rho_k$ for $\mathcal{P}$ is unique up to multiplication of ρ_i by ± 1 , $i = 1 \dots k$.

A Lipschitz field of k -planes on P is approximated by interpolating k -frames on each C_i . Frames in $\mathcal{T}_{n,k}$ are interpolated using $n \times n$ plane rotation matrices [16], $G_i(\theta_i)$, $i = 1 \dots q$, where $q = k(n - k)$. Suppose $e_1 \dots e_n \in \mathbb{R}^n$ are the standard basis vectors. The first $n - k$ matrices are: G_1 is rotation in the plane span $\{e_k, e_n\}$, G_2 in the plane span $\{e_k, e_{n-1}\}$, $\dots$, G_{n-k} in the plane span $\{e_k, e_{k+1}\}$. The second $n - k$ matrices are: G_{n-k+1} is rotation in the plane span $\{e_{k-1}, e_n\}$, G_{n-k+2} in the plane span $\{e_{k-1}, e_{n-1}\}$, $\dots$, $G_{2(n-k)}$ in the plane span $\{e_{k-1}, e_{k+1}\}$, and so on. Finally, the last $n - k$ matrices are: $G_{(k-1)(n-k)+1}$ is rotation in the plane span $\{e_1, e_n\}$, $G_{(k-1)(n-k)+2}$ in the plane span $\{e_1, e_{n-1}\}$, $\dots$, $G_{k(n-k)}$ in the plane span $\{e_1, e_{k+1}\}$. As an example, $G_1(\theta_1)$ is defined as follows. Define the $n \times n$ matrix

$$L(\theta) = [e_k, e_n]R(\theta)[e_k, e_n]^T, \quad R(\theta) = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}$$

and the $n \times n$ orthogonal projection matrix $P = [e_k, e_n][e_k, e_n]^T$. Then $G_1(\theta_1) = (I - P) + L(\theta_1)$.

Denote the $n \times k$ matrix $\Psi(\theta_1 \dots \theta_q) = G_1(\theta_1) \dots G_q(\theta_q)[e_1 \dots e_k]$. The usefulness of the plane rotation matrices G_i is indicated in the following lemma.

Lemma 2. For any k -frame $F \in \mathcal{T}_{n,k}$, there exists $(\theta_1 \dots \theta_q) \in \mathbb{R}^q$ such that $F = \Psi(\theta_1 \dots \theta_q)$.

Proof. Suppose $[\rho_1 \dots \rho_k]$ is a lower triangular orthonormal frame (7). A Givens reduction is used to reduce $[\rho_1 \dots \rho_k]$ to $[e_1 \dots e_k]$. More precisely, $[\rho_1 \dots \rho_k]$ is multiplied on the left by the plane rotation matrices $G_1(\theta_1) \dots G_q(\theta_q)$ defined above, where θ_i , $i = 1 \dots q$, are chosen appropriately [16]. In more detail, $G_{n-k}(\theta_{n-k}) \dots G_1(\theta_1)[\rho_1 \dots \rho_k]$ has zeros in column k below row k for some $\theta_1 \dots \theta_{n-k} \in (-\pi/2, \pi/2]$. The diagonal entry in column k is ± 1 and is the only non-zero entry in column k . By allowing $\theta_{n-k} \in (-\pi, \pi]$ if necessary, the diagonal entry in column k becomes 1. Since the columns are orthonormal, the diagonal entry is also the only non-zero entry in row k . Next, $G_{n-k+1} \dots G_{2(n-k)}$ produce zeros in column $k-1$ below row k . (Row k already has a zero.) The diagonal entry becomes the only non-zero entry in column $k-1$ or row $k-1$. Continuing in this fashion, $G_q(\theta_q) \dots G_1(\theta_1)[\rho_1 \dots \rho_k] = [e_1 \dots e_k]$. Thus, any lower triangular k -frame $[\rho_1 \dots \rho_k]$, satisfies $[\rho_1 \dots \rho_k] = G_1(-\theta_1) \dots G_q(-\theta_q)[e_1 \dots e_k]$, for some $(\theta_1 \dots \theta_q) \in \mathbb{R}^q$. This completes the proof. $\square$

The field μ_1 which approximates $N^s(V)$ is constructed from a given finite set of data points in $\mathcal{G}_{n,k_1}$. These data points are the k_1 -planes $\{N_y^s(V) : y \in \mathcal{C}^0\}$, where $\mathcal{C}^0$ is the set of vertices of $\mathcal{C}$. The field μ_1 on P is formulated by specifying μ_1 on each C_i in a consistent manner. Denote by $y_{1,i} \dots y_{d+1,i} \in \mathbb{R}^n$ the vertices of C_i . According to lemma 1, there is a reduced orthonormal frame $E(i, y_{1,i})$ for $N_{y_{1,i}}^s(V)$. For small H , $N_{y_{2,i}}^s(V) \dots N_{y_{d+1,i}}^s(V)$ are all near $N_{y_{1,i}}^s(V)$. Hence, by the discussion following lemma 1, there is a unique reduced orthonormal frame $E(i, y_{j,i})$ for $N_{y_{j,i}}^s(V)$ near $E(i, y_{1,i})$, $j = 2 \dots d+1$.

Now, lemma 2 is used to interpolate the reduced orthonormal frames $E(i, y_{j,i})$, $j = 1 \dots d+1$, over each C_i , $i = 1 \dots N$. According to lemma 2, $E(i, y_{j,i}) = \Psi(\zeta(i, y_{j,i}))$, some $\zeta(i, y_{j,i}) \in \mathbb{R}^q$, $j = 1 \dots d+1$ for $i = 1 \dots N$. The $\zeta(i, y_{j,i}) \in \mathbb{R}^q$ are interpolated using an affine map $\lambda_i : C_i \rightarrow \mathbb{R}^q$. In fact, there is a unique affine map $\lambda_i : C_i \rightarrow \mathbb{R}^q$ satisfying the conditions $\lambda_i(y_{m,i}) = \zeta(i, y_{m,i})$ for $m = 1 \dots d+1$. The interpolated orthonormal frames are then $\Psi \circ \lambda_i(y) \in \mathcal{T}_{n,k}$ for $y \in C_i$, $i = 1 \dots N$. The field $\mu_1 : P \rightarrow \mathcal{G}_{n,k_1}$ is defined by $\mu_1(y) = \psi \circ \lambda_i(y)$ for $y \in C_i$, $i = 1 \dots N$, where $\psi(\theta_1 \dots \theta_q) = \text{Range}\{\Psi(\theta_1, \dots, \theta_q)\}$.

To address the consistency of the interpolation, suppose the $(d-1)$ -simplex $D = C_i \cap C_l$ is a common face of two d -simplices. For consistency, it is necessary that at a point $x \in D$, the k -plane at x derived from interpolation on C_i agrees with the k -plane at x derived from interpolation on C_l . To guarantee this, the data on D for λ_i must differ from the data on D for λ_l by a certain constant translation. To be precise, suppose $x_j \in \mathbb{R}^n$, $j = 1 \dots d$, are the vertices of D . For each k_1 -plane $N_{x_j}^s(V)$, two reduced orthonormal frames $E(i, x_j)$ and $E(l, x_j)$ of $N_{x_j}^s(V)$ have been defined. Thus, $E(i, x_j) = E(l, x_j)Q(x_j)$, where $Q(x_j)$ is a diagonal matrix with diagonal entries ± 1 . Now, $Q(x_j) = E(i, x_j)^T E(l, x_j)$, and the set $\{E(i, x_j) : j = 1 \dots d\}$ is contained in an arbitrarily small neighbourhood for small H , as is $\{E(l, x_j) : j = 1 \dots d\}$. Hence, $Q(x_j) = Q$ is a constant function of $j = 1 \dots d$. Thus, $E(i, x_j) = E(l, x_j)Q$, $j = 1 \dots d$, where Q is a diagonal matrix with diagonal entries ± 1 . This implies that, if $E(i, x_j) = \Psi(\zeta_j)$ for some $\zeta_j \in \mathbb{R}^q$, $j = 1 \dots d$, then $E(l, x_j) = \Psi(\zeta_j + c)$, $j = 1 \dots d$, where the components c_i of $c \in \mathbb{R}^q$ are

$$c_i = \begin{cases} \pi & \text{if } i = j(n - k_1) \text{ and } Q_{jj} = -1 \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

for $i = 1 \dots q$. Here, Q_{jj} is the diagonal entry (j, j) of Q . In conclusion, $\zeta(i, y_{j,i}) \in \mathbb{R}^q$ may be chosen so that, on a common face $D = C_i \cap C_l$ with vertices $x_j \in \mathbb{R}^n$, $j = 1 \dots d$, if $E(i, x_j)$ and $E(l, x_j)$ are two reduced orthonormal frames of $N_{x_j}^s(V)$, then $\zeta(i, x_j) - \zeta(l, x_j) = c$, where $c \in \mathbb{R}^q$ has form (8).

As a consequence of the above relationship between the data $\zeta(i, x_j)$ and $\zeta(l, x_j)$, $j = 1 \dots d$, the definition $\mu_1(y) = \psi \circ \lambda_i(y)$, $y \in C_i$, $i = 1 \dots N$ is consistent. This is because on a common face of two d -simplices $D = C_i \cap C_l$, with vertices $x_j \in \mathbb{R}^n$, $j = 1 \dots d$, $\psi \circ \lambda_i$ and $\psi \circ \lambda_l$ agree. In fact, $\lambda_i|_D : D \rightarrow \mathbb{R}^q$ is the unique affine map satisfying the conditions $\lambda_i(x_j) = \zeta(i, x_j)$, $j = 1 \dots d$, while $\lambda_l|_D : D \rightarrow \mathbb{R}^q$ is the unique affine map satisfying the conditions $\lambda_l(x_j) = \zeta(l, x_j)$, $j = 1 \dots d$. Since $\zeta(i, x_j) - \zeta(l, x_j) = c$, $j = 1 \dots d$, where c is of the form (8), we have $\lambda_i(x) - \lambda_l(x) = c$ for $x \in D$. This implies that $\Psi \circ \lambda_i(x) = \Psi \circ \lambda_l(x)Q$ for $x \in D$, where Q is a diagonal matrix whose diagonal entries are ± 1 . Thus, $\psi \circ \lambda_i(x) = \psi \circ \lambda_l(x)$, $x \in D$. Therefore, μ_1 is continuous on P .

The construction of μ_2 from the data points $\{N_p^u(V) : p \in C^0\}$ is analogous to the construction of μ_1 . Now, $\mu(x) = \mu_1(x) \oplus \mu_2(x)$ is a k_0 -plane for $x \in P$ provided H is small. This is because hyperbolicity implies that $\{\angle N^s(V), N^u(V)\} > 0$ and μ_1, μ_2 give uniform approximations to $N^s(V)$, $N^u(V)$, respectively, as $H \rightarrow 0$. The field of k_0 -planes $\mu : P \rightarrow \mathcal{G}_{n,k_0}$ induces a vector bundle $N(P)$ with base space P , whose fibre at $x \in P$ is the k_0 -plane $\mu(x)$.

The vector bundle $N(P)$ induces a tubular neighbourhood of P if it is transverse to P and locally Lipschitz [26]. Hence, $N(P)$ induces a tubular neighbourhood due to the following lemma.

Lemma 3. *The vector bundle $N(P)$ is transverse to P and Lipschitz, uniformly for $H \rightarrow 0$.*

Proof. The transversality of $N(P)$ is implied by the following stronger condition. Each d -simplex subspace P_i , $i = 1 \dots N$, of P is a manifold with boundary with tangent bundle $T(P_i)$. There exists $\alpha > 0$ such that

$$\inf\{\angle N_x(P), T_x(P_i) : \text{all } P_i \text{ containing } x, x \in P\} \geq \alpha \quad (9)$$

for $H \rightarrow 0$. This property is inherited from $N(V)$, $T(V)$. To be specific, because the given $N(V)$ is transverse and V is compact, we have $\{\angle N(V), T(V)\} > 0$. Since $N(P)$ gives a uniform approximation to $N(V)$ and P approximates V in the Lipschitz norm as $H \rightarrow 0$, (9) follows.

Like transversality, the uniform Lipschitz property of $N(P)$ is also inherited from $N(V)$. For simplicity, consider the attracting case. Recall that $\mu_1 : P \rightarrow \mathcal{G}_{n,k_1}$ is $\mu_1(y) = \psi \circ \lambda_i(y)$ for $y \in C_i$, $i = 1 \dots N$. The data for λ_i are $\lambda_i(y_{m,i}) = \zeta(i, y_{m,i})$ for $m = 1 \dots d+1$, where $y_{1,i} \dots y_{d+1,i} \in \mathbb{R}^n$ are the vertices of C_i . Because λ_i is an affine interpolant, $\text{Lip}\{\lambda_i\} \leq (\text{constant}/H) \cdot \max\{|\zeta(i, y_{j,i}) - \zeta(i, y_{m,i})| : j, m = 1 \dots d+1\}$ for $H \rightarrow 0$, [9]. Since $N(V)$ is Lipschitz, $\max\{|\zeta(i, y_{j,i}) - \zeta(i, y_{m,i})| : j, m = 1 \dots d+1\} \leq \text{constant} \cdot \max\{|y_{j,i} - y_{m,i}| : j, m = 1 \dots d+1\}$. Combining these, we obtain $\text{Lip}\{\lambda_i\} \leq \text{constant}$ uniformly for $H \rightarrow 0$, $i = 1 \dots N$. This implies that μ_1 is Lipschitz, uniformly in $H \rightarrow 0$. $\square$

3.4. Sections of the transverse bundle

The field of k_0 -planes $\mu : P \rightarrow \mathcal{G}_{n,k_0}$ induces a vector bundle $N(P)$ with base space P , whose fibre at $x \in P$ is the k_0 -plane $\mu(x)$. This $N(P)$ gives a tubular neighbourhood of P . Analogous to the approach in section 2, we work in a neighbourhood of the zero section in $N(P)$, which is equivalent to a neighbourhood of P in $\mathbb{R}^n$. Any C^r , $r \geq 1$, manifold $\tilde{V}$ Lipschitz-near V corresponds to the graph of a section σ of $N(P)$, for small H . The section σ is C^r on each C_i . A candidate manifold $\tilde{V}$ is approximated by a section σ_D of $N(P)$ which is polynomial on each C_i in appropriate coordinates. On each C_i , σ_D is a polynomial map into the fibres of $N(P)$. In fact, $N(P) = N^s(P) \oplus N^u(P)$, where the fibre of $N^s(P)$ at $x \in P$ is

Figure 3. Approximation to $\tilde{V}$, attracting case, $p = 2$.

the k_1 -plane $\mu_1(x)$ and the fibre of $N^u(P)$ at $x \in P$ is the k_2 -plane $\mu_2(x)$. The approximating section is $\sigma_D(x) = (x, v^s(x), v^u(x))$, where $v^s(x) \in N_x^s(P)$, $v^u(x) \in N_x^u(P)$. In appropriate coordinates, on each C_i , v^s and v^u are Lagrange polynomials of order $p \geq 1$, [7].

The section σ on C_i is approximated by interpolating a discrete data set consisting of the values of σ at certain points of C_i . In this way, the locally C^r section σ is approximated by a discrete section σ_D . The discrete data set for σ_D on C_i consists of the points of intersection of the graph of σ in $N_{C_i}(P)$ with the fibres $N_x(P)$, for points x in the principal lattice of order p of C_i . See figure 3. The principal lattice of order p of C_i , denoted Σ_i , is the set of points in C_i with barycentric coordinates $b_1 \dots b_{d+1} \in \{0, 1/p \dots (p-1)/p, 1\}$, [7]. Denote the points of Σ_i by $x_{i,j} \in C_i \subset P$, $j = 1 \dots m$. Then the points of intersection of the graph of σ in $N_{C_i}(P)$ with the fibres $N_x(P)$, $x \in \Sigma_i$, are

$$(x_{i,j}, v_{i,j}^s, v_{i,j}^u) \in N_{C_i}(P), \quad \text{for some } v_{i,j}^s \in N_{x_{i,j}}^s(P), v_{i,j}^u \in N_{x_{i,j}}^u(P),$$

$j = 1 \dots m$. The discrete section σ_D is composed of stable and unstable parts, $v^s(x)$ and $v^u(x)$. Here, $v^s(x)$, $x \in C_i$, is fitted to $v_{i,j}^s$, $j = 1 \dots m$, and $v^u(x)$, $x \in C_i$, is fitted to $v_{i,j}^u$, $j = 1 \dots m$.

Before constructing $v^s(x)$, we compare the present setting with a typical application of the Lagrange polynomials. The usual application gives an interpolating Lagrange polynomial $C_i \rightarrow \mathbb{R}^{k_0}$, [7]. This polynomial serves to approximate a graph in $C_i \times \mathbb{R}^{k_0}$. In the present case, the interpolant is a section $C_i \rightarrow N_{C_i}(P)$ which approximates a graph in $N_{C_i}(P)$. Thus, to proceed, coordinates on $N_{C_i}^s(P)$, $i = 1 \dots N$, are induced by smooth orthonormal moving frames. Namely, an orthonormal basis of $N_x^s(P)$ is given by the columns of the $n \times k_1$ matrix $E_i(x) = \Psi \circ \lambda_i(x)$ for $x \in C_i$. For each $x \in C_i$, this matrix induces an invertible linear transformation $\xi_i(x) : \mathbb{R}^{k_1} \rightarrow N_x^s(P)$, $\xi_i(x)(\rho) = E_i(x)\rho$. There is a unique Lagrange polynomial $\eta_i^s : C_i \rightarrow \mathbb{R}^{k_1}$ of total degree p fitting the data

$$\eta_i^s(x_{i,j}) = \xi_i(x_{i,j})^{-1}(v_{i,j}^s), \quad j = 1 \dots m,$$

[7, 9]. Now put $v^s(x) = \xi_i(x) \circ \eta_i^s(x)$ for $x \in C_i$. This definition is consistent because on a common $(d-1)$ -face D of two d -simplices C_i and C_l , $\xi_i(x) \circ \eta_i^s(x)$ and $\xi_l(x) \circ \eta_l^s(x)$ agree. To prove this, it is sufficient to show $\eta_i^s(x) = \xi_i(x)^{-1} \circ \xi_l(x) \circ \eta_l^s(x)$ for $x \in D$. The equality holds for those $x = x_{i,j}$, $j = 1 \dots m$, contained in D (that is, on the principal lattice of order p of D). Now, $\xi_i(x)^{-1} \circ \xi_l(x) = Q$ for $x \in D$, where Q is a diagonal matrix whose diagonal entries are ± 1 . In particular, $\xi_i(x)^{-1} \circ \xi_l(x)$, $x \in D$, is constant. Hence, $\eta_i^s(x)$ and $\xi_i(x)^{-1} \circ \xi_l(x) \circ \eta_l^s(x)$, $x \in D$, are both polynomials of total degree p fitting the same data on the principal lattice of order p of D , proving the equality on D .

The construction of v^u is analogous to the construction of v^s . The resulting approximating section $\sigma_D(x) = (x, v^s(x), v^u(x))$ of $N(P)$ is continuous. If $\tilde{V}$ is of smoothness class C^{p+1} , σ_D is an approximation to σ of order p . That is, $\sup\{|v(x) - v_D(x)|_x : x \in P\} = O(H^{p+1})$ as $H \rightarrow 0$, where $\sigma(x) = (x, v(x))$ and $\sigma_D(x) = (x, v_D(x))$.

Remarks

- (i) In the case that V is a flat manifold, the method of discretizing manifolds given in section 3 reduces to the C^0 finite element method. The finite element method is typically used to approximate manifolds which are globally graphs over a flat space [7]. The approach of section 3 generalizes this to include other manifolds. Moreover, it does this while keeping the dimension of the domain and the range of the finite element basis functions [7, 9] the same as in the usual case of a graph over a flat space.

4. The discrete graph transform

In this section the continuation algorithm sketched in section 2 is formulated in a way which is realizable in a numerical implementation. The approximation of candidate manifolds described in section 3 leads to a discrete version of the graph transform for approximating the perturbed manifold $\tilde{V} \subset \mathbb{R}^n$ in sections 4.1 and 4.2. Similarly, the approximation of the (un)stable bundle leads to a discrete version of the linear graph transform for approximating the hyperbolic splitting of $\tilde{V}$ in section 4.3.

The discrete graph transform Γ_D is formulated by replacing the components of the graph transform described in section 2 with the discrete counterparts of section 3. Namely, $N(V) = N^u(V) \oplus N^s(V)$ is replaced by $N(P) = N^u(P) \oplus N^s(P)$ in section 4.1 and the sections σ of $N(V)$ are replaced by discrete sections σ_D of $N(P)$ in section 4.2. Then, for example, in the attracting case, the graph transform Γ of σ_D is the section $\Gamma(\sigma_D)$ whose graph is equal to the $\tilde{F}$ -image of the graph of σ_D . The section $\Gamma(\sigma_D)$ is not a discrete section. So the output of Γ_D matches its input, $\Gamma(\sigma_D)$ is approximated by a discrete section $\mathcal{I} \circ \Gamma(\sigma_D) = \Gamma_D(\sigma_D)$. For the discrete graph transform Γ_D to be well behaved, it must preserve the C^0 norm and the Lipschitz constant of sections. In section 4.2 it will be seen that these properties of Γ_D can be obtained by varying a parameter of the graph transform.

4.1. The graph transform of sections of $N(P)$

In this section, we follow section 2.2, replacing $N(V) = N^u(V) \oplus N^s(V)$ by $N(P) = N^u(P) \oplus N^s(P)$. The graph transform formulated in terms of $N(P)$ may be considered a special case of the usual general description of the graph transform. The difference between this section and section 2.2 is that here $N(P)$ is Lipschitz rather than smooth. To be specific, the Lipschitz constant of a section σ^s of $N^s(P)$ is defined as follows. First, $N(P)$ induces a homeomorphism $\psi : V \rightarrow P$. Suppose $N^s(V)$ is the vector bundle over V whose fibre at $p \in V$ is $N_{\psi(p)}^s(P)$. Since $N^s(V)$ is a subbundle of $T_V(M)$, the Lipschitz constant of the section $\sigma^s \circ \psi$ of $N^s(V)$ is defined in section 2. Hence, $\text{Lip}\{\sigma^s\} = \text{Lip}\{\sigma^s \circ \psi\}$, and similarly for σ^u . Now, $\text{Lip}\{\sigma\}$ for a section σ of $N(P)$ is defined as in section 2.2. Suppose $Z = Z(\epsilon) = \{(x, v) \in N(P) : |v|_x \leq \epsilon\}$ and $\mathcal{S}_{\epsilon, \delta} = \{\sigma : P \rightarrow Z : \text{Lip}(\sigma) \leq \delta\}$. The space $\mathcal{S}_{\epsilon, \delta}$ with the C^0 norm $\|\cdot\|$ described in section 2.2 is complete.

Given $\sigma \in \mathcal{S}_{\epsilon, \delta}$, $\sigma(x) = (x, v^s(x), v^u(x))$, the graph transform of σ is a section $\Gamma(\sigma)(x) = (x, w^s(x), w^u(x))$ of $N(P)$. Here, $w^s(x)$ is the stable part of the intersection of the $\tilde{F}^0$ -image of the graph of σ with the fibre $N_x(P)$. Thus, to define $w^s(x)$ for a given $x \in P$, first solve

$$x = \pi \circ \tilde{F}^0(p, v^s(p), v^u(p)), \quad (10)$$

for $p \in P$, where $\pi : N(P) \rightarrow P$ is the vector bundle projection. The $\tilde{F}^0$ -image of the graph of σ transversally intersects the fibre $Z_x(P)$. In (10) we solve for the unique $p \in P$ such that

$\tilde{F}^0 \circ \sigma(p)$ is contained in the fibre $Z_x(P)$. Equation (10) has a unique solution for $p \in P$, provided ϵ, δ, θ and H are small. Denote this solution by $p = p(x, v^s, v^u)$. Now, $w^s(x)$ is given by the formula

$$w^s(x) = P_x^s \circ \tilde{F}^0(p, v^s(p), v^u(p)), \quad (11)$$

for $x \in P$, where $P_x^s : N_x(P) \rightarrow N_x(P)$ is the linear projection with range $N_x^s(P)$ and nullspace $N_x^u(P)$.

The unstable part w^u is defined implicitly by eliminating x in

$$v^u(x) = P_x^u \circ \tilde{F}^0(p, v^s(p), w^u(p)), x = \pi \circ \tilde{F}^0(p, v^s(p), w^u(p)), \quad (12)$$

for $p \in P$, where $P_x^u : N_x(P) \rightarrow N_x(P)$ is the linear projection with range $N_x^u(P)$ and nullspace $N_x^s(P)$. In words, w^u is the unstable part supplementing the stable part of σ such that the $\tilde{F}^0$ -image of the graph of the resulting section agrees with σ in its unstable part. In (12) we solve for the vector $w = w^u(p) \in Z_p^u(P)$ such that the $\tilde{F}^0$ -image of $(p, v^s(p), w)$ has an unstable component in the graph of v^u . There is a unique solution for $w^u(p)$ in (12) for small ϵ, δ, θ and H . In both (10) and (12), there are unique solutions since P is Lipschitz-near V and $N(P)$ is a uniform approximation to $N(V)$ as $H \rightarrow 0$. The proof uses the Lipschitz implicit function theorem [14, p 207]. As in section 2.2, by replacing $\tilde{F}$ with $\tilde{F}^N$ if necessary, Γ becomes a contraction on $\mathcal{S}_{\epsilon, \delta}$ whose fixed point gives the $\tilde{F}$ -invariant manifold $\tilde{V}$.

4.2. The discrete graph transform

In this section, the formulation of Γ_D started in section 4.1 is finished. In a computational setting, only discrete sections may be used. Thus, the domain of Γ from section 4.1 is restricted to the subset of $\mathcal{S}_{\epsilon, \delta}$ consisting of discrete sections. For $\sigma_D \in \mathcal{S}_{\epsilon, \delta}$, where σ_D is a discrete section of the form constructed in section 3.4, $\Gamma(\sigma_D)$ is not a discrete section. To iterate the discrete graph transform, $\Gamma_D(\sigma_D)$ must be a discrete section. Thus, define $\Gamma_D(\sigma_D) = \mathcal{I} \circ \Gamma(\sigma_D)$, where $\mathcal{I} \circ \sigma$ is the discrete section approximating σ described in section 3.4. Whether Γ_D leaves $\mathcal{S}_{\epsilon, \delta}$ invariant depends on the effect $\mathcal{I}$ has on both the C^0 norm and the Lipschitz constant of sections in $\mathcal{S}_{\epsilon, \delta}$. This calls for a more detailed analysis of $\mathcal{I}$.

To analyse the effect $\mathcal{I}$ has on the C^0 norm and the Lipschitz constant of sections, a formula for $\mathcal{I}(\sigma)$ is obtained. A section $\sigma \in \mathcal{S}_{\epsilon, \delta}$ is

$$\sigma(x) = (x, \xi_i^s(x) \circ f_i^s(x), \xi_i^u(x) \circ f_i^u(x)), \quad x \in \mathcal{C}_i \quad (13)$$

for some $f_i^s : \mathcal{C}_i \rightarrow \mathbb{R}^{k_1}$ and $f_i^u : \mathcal{C}_i \rightarrow \mathbb{R}^{k_2}$. Here, ξ_i^s and ξ_i^u are defined in section 3.4. Recall that $\xi_i^s(x) : \mathbb{R}^{k_1} \rightarrow N_x^s(P)$, $\xi_i^s(x)(\rho) = E_i^s(x)\rho$, where the columns of the $n \times k_1$ matrix $E_i^s(x)$ form an orthonormal basis for $N_x^s(P)$, $x \in \mathcal{C}_i$. The description of $\xi_i^u(x)$ is analogous. Recall that Σ_i , defined in section 3.4, is the principal lattice of order $p \geq 1$ of the d -simplex $\mathcal{C}_i$. Then $\mathcal{I}(\sigma)$ is the discrete section σ_D of $N(P)$ whose data on $\mathcal{C}_i$ consist of the points of intersection of the graph of σ in $N_{\mathcal{C}_i}(P)$ with the fibres $N_x(P)$, $x \in \Sigma_i$. To be specific,

$$\mathcal{I}(\sigma)(x) = (x, \xi_i^s(x) \circ L_i^s \circ f_i^s(x), \xi_i^u(x) \circ L_i^u \circ f_i^u(x))$$

for $x \in \mathcal{C}_i$, where L_i^s and L_i^u are the standard Lagrange interpolation operators on functions on $\mathcal{C}_i$. Here, the Lagrange interpolation operators are defined as follows. Given $f : \mathcal{C}_i \rightarrow \mathbb{R}^{k_1}$, $L_i^s \circ f : \mathcal{C}_i \rightarrow \mathbb{R}^{k_1}$ is the unique polynomial of total degree p with $L_i^s \circ f(x) = f(x)$ for $x \in \Sigma_i$. The definition of L_i^u is analogous.

The maximum factor of growth of the C^0 norm of a section under $\mathcal{I}$ is $C_p = \sup\{\|\mathcal{I}(\sigma)\|/\|\sigma\| : \sigma \in \mathcal{S}_{\epsilon, \delta}\}$. The maximum factor of growth of the Lipschitz constant of a section under $\mathcal{I}$ is $C'_p = \sup\{\text{Lip}\{\mathcal{I}(\sigma)\}/\text{Lip}\{\sigma\} : \sigma \in \mathcal{S}_{\epsilon, \delta}\}$. Here, C_p and C'_p are bounded as $H \rightarrow 0$. Since L_i^s and L_i^u are linear operators and $\xi_i^s(x)$ and $\xi_i^u(x)$ are isometries, the Lipschitz

constant of $\mathcal{I}$ is also bounded by C_p for $p \geq 1$. Thus, ideally, in the case $C_p = C'_p = 1$, $\mathcal{I}$ has no deleterious effect on Γ , and Γ_D is a contraction on $\mathcal{S}_{\epsilon, \delta}$ with no adjustments to any parameters.

Now, estimates of the factors of growth C_p and C'_p are obtained in the case $p = 1$. Consider the effect of $\mathcal{I}$ on a section $\sigma \in \mathcal{S}_{\epsilon, \delta}$ of the form (13). Here, $\text{graph}\{f_i^s\}$ is contained in the convex set $\mathcal{C}_i \times \{v \in \mathbb{R}^{k_1} : |v| \leq \epsilon\}$ and $\text{graph}\{L_i^s \circ f_i^s\}$ is contained in the convex hull of a set of points in $\text{graph}\{f_i^s\}$. Thus,

$$\text{graph}\{L_i^s \circ f_i^s\} \subset \mathcal{C}_i \times \{v \in \mathbb{R}^{k_1} : |v| \leq \epsilon\}$$

and hence $|\mathcal{I}^s(\sigma)|_s \leq \epsilon$, where $\mathcal{I}^s(\sigma)(x) = (x, \xi_i^s(x) \circ L_i^s \circ f_i^s(x))$. Similarly, $|\mathcal{I}^u(\sigma)|_u \leq \epsilon$ and thus $\|\mathcal{I}(\sigma)\| \leq \epsilon$. In other words, $C_p = 1$ for $p = 1$. Now, $\text{Lip}\{\mathcal{I}(\sigma)\}$ may be large if any interior angles of the $\mathcal{C}_i$ are small. In practice, it is possible to keep the interior angles of the $\mathcal{C}_i$ nearby those of the standard d -simplex D . A calculation shows that, in the case that each $\mathcal{C}_i$ is the image of D under an affine conformal mapping, $\text{Lip}\{\mathcal{I}(\sigma)\} \leq \sqrt{d} \text{Lip}\{\sigma\}$. In conclusion, $C_p = 1$ and $C'_p = \sqrt{d}$ for $p = 1$.

Difficulties caused by either $C_p > 1$ or $C'_p > 1$ may be offset by a modification of one of the parameters of Γ . For simplicity, consider the attracting case. Suppose that $0 < \alpha < 1$ is the factor of (weakest) normal contraction towards V under F . Also, $0 < \mu < 1$ from (2) is a bound on $\alpha/\{\text{the factor of (strongest) tangential contraction under } F\}$. Given $\sigma \in \mathcal{S}_{\epsilon, \delta}$, the C^0 norm and Lipschitz constant of $\Gamma(\sigma)$ are multiplied by factors $c\alpha^N + o(1)$ and $c\mu^N + o(1)$, respectively, as $\epsilon + \delta + \theta + H \rightarrow 0$. The C^0 norm and Lipschitz constant of $\Gamma_D(\sigma)$ are multiplied by factors $C_p c\alpha^N + o(1)$ and $C'_p c\mu^N + o(1)$, respectively. Thus, by choosing N large enough, we obtain $\Gamma_D : \mathcal{S}_{\epsilon, \delta} \rightarrow \mathcal{S}_{\epsilon, \delta}$. Also, Γ_D is a contraction since

$$\text{Lip}\{\Gamma_D\} \leq \text{Lip}\{\mathcal{I}\}\text{Lip}\{\Gamma\} = C_p c\alpha^N + o(1)$$

as $\epsilon + \delta + \theta + H \rightarrow 0$.

Alternatively, it is possible to estimate $\text{Lip}\{\mathcal{I}(\sigma)\}$ using the constant $C''_p = H \sup\{\text{Lip}\{\mathcal{I}(\sigma)\}/\|\sigma\| : \sigma \in \mathcal{S}_{\epsilon, \delta}\}$, which is bounded as $H \rightarrow 0$. In this case, there exists a constant $c > 0$ and a positive function $\omega(H) \rightarrow 0$ as $H \rightarrow 0$, such that the following holds. If $\epsilon = cH\delta$, $\omega(H) < c\delta$, $\theta < c\epsilon$, δ is sufficiently small and N sufficiently large, then $\Gamma_D : \mathcal{S}_{\epsilon, \delta} \rightarrow \mathcal{S}_{\epsilon, \delta}$ is a contraction [2]. This result does not use the full hypothesis of normal hyperbolicity, but only the existence of a C^1 , 0-normally hyperbolic manifold $\tilde{V}$, [27]. This explains why Γ_D is a contraction, in practice, for some dynamical systems even in the absence of normal hyperbolicity. This result may also indicate why the algorithm converges in some cases for large H , since this bound on $\text{Lip}\{\mathcal{I}(\sigma)\}$ decreases as H increases.

In either of the scenarios in the preceding two paragraphs, Γ_D has a fixed point $\sigma_D^* \in \mathcal{S}_{\epsilon, \delta}$ whose graph approximates $\tilde{V}$. In conclusion, the following result may be stated.

Theorem 1. *Suppose there exists an $\tilde{F}$ -invariant manifold $\tilde{V}$ Lipschitz-near a given F -invariant manifold V , both manifolds C^r , $r \geq 1$. In addition, suppose V and $\tilde{V}$ are ρ -normally hyperbolic, $\rho \geq 0$, where the splitting of $N(V)$ is given. Then for ϵ, δ, θ and H small and N sufficiently large*

- (i) *the discrete graph transform $\Gamma_D : \mathcal{S}_{\epsilon, \delta} \rightarrow \mathcal{S}_{\epsilon, \delta}$ is a contraction,*
- (ii) *if $r \geq 2$ or the order of approximation $p = 1$, then the fixed point $\sigma_D^* \in \mathcal{S}_{\epsilon, \delta}$ of Γ_D satisfies $\phi \circ \sigma_D^*(P) \rightarrow \tilde{V}$ in the Lipschitz norm as $H \rightarrow 0$,*
- (iii) *if $r \geq p + 1$, then $\phi \circ \sigma_D^*(P)$ is a C^0 approximation to $\tilde{V}$ of order p .*

As discussed above, the estimates used in the proof differ between the 0-normally hyperbolic case and the ρ -normally hyperbolic case, $\rho \geq 1$. This means that in these two cases, there is a different required relationship among the parameters $\epsilon, \delta, \theta, H$ and N . In addition, although

the hypotheses of theorem 1 allow the case that $\tilde{V}$ is 0-normally hyperbolic, generally the existence of a C^r invariant manifold $\tilde{V}$ is not known without the hypothesis of r -normally hyperbolicity, $r \geq 1$, [27].

Remarks

- (i) In section 4.2, estimates on the sizes of C_p, C'_p were obtained which hold in general. However, in practice, the minimal sizes of C_p, C'_p for a specific computation may depend on the shape of the manifold $\tilde{V}$.
- (ii) In practice, it is possible to use a cut-off function to control $\|\mathcal{I}(\sigma)\|$, making $C_p = 1$ for $p \geq 1$. Thus, the control of $\text{Lip}\{\mathcal{I}(\sigma)\}$ becomes an important numerical issue.

4.3. The discrete linear graph transform

The next topic is the computation of the approximate hyperbolic splitting of $\tilde{V}$. In section 4.2, an approximation $\phi \circ \sigma_D^*(P)$ to $\tilde{V}$ was obtained for $H \rightarrow 0$. The simplicial complex $\tilde{C}$ with vertices $\phi \circ \sigma_D^*(C^0)$, where C^0 is the set of vertices of P , supports the manifold $\phi \circ \sigma_D^*(P)$. Suppose $\tilde{P} \subset \mathbb{R}^n$ is the polyhedron of $\tilde{C}$ and $N(\tilde{P})$ is a given transverse bundle. Given such an $N(\tilde{P})$, the approximate hyperbolic splitting of $\tilde{V}$ is given by a splitting $N(\tilde{P}) = N^u(\tilde{P}) \oplus N^s(\tilde{P})$.

In this section, the discrete linear graph transforms $\mathcal{L}_D^u$ and $\mathcal{L}_D^s$ are used to determine $N^u(\tilde{P})$ and $N^s(\tilde{P})$. Here, for simplicity, it is assumed that $N(\tilde{P})$ and $N(P)$ are approximately normal in the following sense. Each d -simplex subspace $P_i, i = 1 \dots N$, of P is a manifold with boundary with tangent bundle $T(P_i)$. Then

$$\inf\{\angle N_x(P), T_x(P_i) : \text{all } P_i \text{ containing } x, x \in P\} \rightarrow \pi/2$$

as $H \rightarrow 0$. Next, $\mathcal{L}_D^u$ is formulated. The formulation of $\mathcal{L}_D^s$ is analogous.

The initial data for $\mathcal{L}_D^u$ are a splitting $N(\tilde{P}) = N^{u,0}(\tilde{P}) \oplus N^{s,0}(\tilde{P})$. This splitting is obtained from $N(P) = N^u(P) \oplus N^s(P)$ by parallel translation followed by projection onto the fibres of $N(\tilde{P}) \subset T_{\tilde{P}}(\mathbb{R}^n)$ using Q , as in section 2.3. To be specific, suppose π is the vector bundle projection of $N(P)$. Then $N_y^{u,1}(\tilde{P}), N_y^{s,1}(\tilde{P})$ are obtained from $N_p^u(P), N_p^s(P)$, $p = \pi \circ \phi^{-1}(y)$, by parallel translation $T_p(\mathbb{R}^n) \rightarrow T_y(\mathbb{R}^n)$ along the ϕ -images of fibres of $N(P)$. In the present case, parallel translation is trivially defined by the identification of $T_x(\mathbb{R}^n), x \in \mathbb{R}^n$, with the ambient space $\mathbb{R}^n$. Namely, parallel translation $T_p(\mathbb{R}^n) \rightarrow T_y(\mathbb{R}^n)$, $p = \pi \circ \phi^{-1}(y) \in P, y \in \tilde{P}$, is $v \rightarrow v$. In the present setting,

$$Q : T_{\tilde{P}}(\mathbb{R}^n) \rightarrow N(\tilde{P}) \subset T_{\tilde{P}}(\mathbb{R}^n)$$

is, on each fibre $T_x(\mathbb{R}^n)$, the linear orthogonal projection with range $N_x(\tilde{P})$. The initial data are then

$$N^{u,0}(\tilde{P}) = Q(N^{u,1}(\tilde{P})), N^{s,0}(\tilde{P}) = Q(N^{s,1}(\tilde{P})).$$

This procedure produces non-degenerate initial data for $\epsilon + \delta + \theta + H \rightarrow 0$.

As in section 2.3, $L(\tilde{P})$ is the bundle whose fibre at $y \in \tilde{P}$ is the space of linear transformations $N_y^{u,0}(\tilde{P}) \rightarrow N_y^{s,0}(\tilde{P})$. The domain of $\mathcal{L}_D^u$ is a subset of the space of sections $\mathcal{S}_\eta = \{\sigma : \tilde{P} \rightarrow L(\tilde{P}) : \sup_y \|\sigma(y)\| \leq \eta\}$, where the operator norm $\|\cdot\|$ is associated with the Riemann structure on $T_{\tilde{P}}(\mathbb{R}^n)$. The space $\mathcal{S}_\eta$ is complete with respect to the norm $|\sigma| = \sup_y \|\sigma(y)\|$.

The domain of $\mathcal{L}_D^u$ is the subset of $\mathcal{S}_\eta$ consisting of discrete sections. A discrete section in $\mathcal{S}_\eta$ is constructed using the construction of a discrete field of k_2 -planes $\mu : \tilde{P} \rightarrow \mathcal{G}_{n,k_2}$

in sections 3.2 and 3.3. In fact, the space of sections $\mathcal{S}_\eta$ was introduced as a way to talk about k_2 -plane bundles near $N^{u,0}(\tilde{P})$ in $T_{\tilde{P}}(\mathbb{R}^n)$. The k_2 -plane bundle determined by a section $\sigma \in \mathcal{S}_\eta$ is the bundle over $\tilde{P}$ whose fibre at $x \in \tilde{P}$ is $\text{graph}\{\sigma(x)\} \subset N(\tilde{P}) \subset T_{\tilde{P}}(\mathbb{R}^n)$. Similarly, the unique section $\sigma \in \mathcal{S}_\eta$ determined by a given k_2 -plane bundle $N^{u,*}(\tilde{P})$ near $N^{u,0}(\tilde{P})$ is the $\sigma : \tilde{P} \rightarrow L(\tilde{P})$ where $\sigma(x) \in L_x(\tilde{P})$ is the unique element of $L_x(\tilde{P})$ with the $\text{graph}\{\sigma(x)\} = N_x^{u,*}(\tilde{P}) \subset N_x(\tilde{P}) \subset T_x(\mathbb{R}^n)$.

A discrete section σ_D of $L(\tilde{P})$ is constructed from some given data $\{\sigma_D(x) \in L_x(\tilde{P}) : x \in \tilde{C}^0\}$, where $\tilde{C}^0$ is the set of vertices of $\tilde{P}$, as follows. Using the method of section 3.3, construct the field $\mu : \tilde{P} \rightarrow \mathcal{G}_{n,k_2}$ of k_2 -planes determined by the set of k_2 -plane data points

$$\{\text{graph}\{\sigma_D(x)\} \subset N(\tilde{P}) \subset T_{\tilde{P}}(\mathbb{R}^n) : x \in \tilde{C}^0\}.$$

The discrete section σ_D is then uniquely characterized by $\text{graph}\{\sigma_D(x)\} = \mu(x)$, $x \in \tilde{P}$.

As for the discrete graph transform in sections 4.1 and 4.2, the discrete linear graph transform is formulated in two steps. Roughly speaking, the first step discretizes the linear graph transform and the second step discretizes its domain. In the first step, the linear graph transform $\mathcal{L}^u$ is formulated in the present setting, replacing $N(\tilde{V})$ by $N(\tilde{P})$ in section 2.3. This means that instead of a smooth manifold and transverse bundle, here they are only Lipschitz. In addition, the formulation of $\mathcal{L}^u$ in this section is slightly different from the formulation of $\mathcal{L}^u$ in section 2.3 because $\tilde{P}$ is not $\tilde{F}$ -invariant. In the second step, so that the output of $\mathcal{L}_D^u$ matches its input, define $\mathcal{L}_D^u(\sigma_D) = \mathcal{I} \circ \mathcal{L}^u(\sigma_D)$, $\sigma_D \in \mathcal{S}_\eta$. Here, for $\sigma \in \mathcal{S}_\eta$, $\mathcal{I}(\sigma)$ is the discrete section of $L(\tilde{P})$ defined by the data $\{\sigma(x) : x \in \tilde{C}^0\}$.

To formulate $\mathcal{L}^u$, first the invariance condition is derived. To define the mapping Φ induced by $D\tilde{F}$ on $N(\tilde{P})$, suppose π is the vector bundle projection of $N(\tilde{P})$ and $\phi : Z \rightarrow U$ is the homeomorphism, defined in section 2.1, associated with the tubular neighbourhood of $\tilde{P}$ induced by $N(\tilde{P})$. Then the linear mapping induced by $D\tilde{F}_x : T_x(\mathbb{R}^n) \rightarrow T_y(\mathbb{R}^n)$, $y = \tilde{F}(x)$, $x \in \tilde{P}$, on $N(\tilde{P})$ is

$$\Phi = Q \circ \gamma \circ D\tilde{F}|_{N(\tilde{P})} : N(\tilde{P}) \rightarrow N(\tilde{P}).$$

Here $\gamma : T_y(\mathbb{R}^n) \rightarrow T_p(\mathbb{R}^n)$, $p = \pi \circ \phi^{-1}(y)$, $y \in U$, is parallel translation. Note that $y \in U$ for small H because $\tilde{P} \rightarrow \tilde{V}$ in the C^0 norm as $H \rightarrow 0$.

The linear graph transform $\mathcal{L}^u(\sigma)$ of $\sigma \in \mathcal{S}_\eta$ is characterized by the condition $\Phi(\text{graph}\{\sigma(x)\}) = \text{graph}\{\mathcal{L}^u(\sigma)(y)\}$ where $y = \pi \circ \phi^{-1} \circ \tilde{F}(x)$. To calculate $\mathcal{L}^u(\sigma)(y)$ for a given $y \in \tilde{P}$, first solve $y = \pi \circ \phi^{-1} \circ \tilde{F}(x)$ for $x \in \tilde{P}$. Next, given an orthonormal basis $e_1 \dots e_{k_2}$ for $N_y^{u,0}(\tilde{P})$, solve $e_i = P_y^u \circ \Phi(\rho_i, \sigma(x)(\rho_i))$ for $\rho_i \in N_x^{u,0}(\tilde{P})$, $i = 1 \dots k_2$. Then $\mathcal{L}^u(\sigma)(y)$ is given by the formula

$$\mathcal{L}^u(\sigma)(y)(e_i) = P_y^s \circ \Phi(\rho_i, \sigma(x)(\rho_i)),$$

$i = 1 \dots k_2$. If Φ is replaced by Φ^N , then $\mathcal{L}^u : \mathcal{S}_\eta \rightarrow \mathcal{S}_\eta$ is a contraction for $\epsilon + \delta + \theta + \eta + H$ small and N large.

It remains to verify $\mathcal{L}_D^u(\sigma_D) \in \mathcal{S}_\eta$ for $\sigma_D \in \mathcal{S}_\eta$ and that $\mathcal{L}_D^u : \mathcal{S}_\eta \rightarrow \mathcal{S}_\eta$ is a contraction. Recall $\mathcal{L}_D^u(\sigma_D) = \mathcal{I} \circ \mathcal{L}^u(\sigma_D)$ for $\sigma_D \in \mathcal{S}_\eta$. Thus, to do this, the norm of $\mathcal{I}(\sigma)$, $\sigma \in \mathcal{S}_\eta$ and the Lipschitz constant of $\mathcal{I}$ on $\mathcal{S}_\eta$ are estimated. For $\sigma \in \mathcal{S}_\eta$, $|\mathcal{I}(\sigma)| \leq \eta + o(1)$ and $\text{Lip}\{\mathcal{I}\} = 1 + o(1)$ as $H \rightarrow 0$. Thus, $\mathcal{L}_D^u : \mathcal{S}_\eta \rightarrow \mathcal{S}_\eta$ is a contraction for $\epsilon + \delta + \theta + \eta + H$ small and N large.

The fixed point $\sigma_D^* \in \mathcal{S}_\eta$ of $\mathcal{L}_D^u$ gives an approximation to $N^u(\tilde{V})$. In conclusion, the following result may be stated.

Theorem 2. *Suppose there is a given $\tilde{F}$ -invariant manifold $\tilde{V}$ Lipschitz-near a given F -invariant manifold V , both manifolds C^r , $r \geq 1$. Suppose V and $\tilde{V}$ are ρ -normally*

hyperbolic, $\rho \geq 0$, where the splitting of $N(V)$ is given. Then, for $\epsilon, \delta, \theta, \eta$ and H small and N sufficiently large:

- (i) the discrete linear graph transform $\mathcal{L}_D^u : \mathcal{S}_\eta \rightarrow \mathcal{S}_\eta$ is a contraction,
- (ii) the fixed point $\sigma_D^* \in \mathcal{S}_\eta$ of $\mathcal{L}_D^u$ gives an approximation to $N^u(\tilde{V})$ in the following sense.
 Suppose $\gamma : N_x(\tilde{V}) \rightarrow N_y(\tilde{P})$, $y = \pi \circ \phi^{-1}(x)$, is parallel translation and σ is a section of $L(\tilde{P})$ satisfying $\text{graph}\{\sigma(y)\} = \gamma(N_x^u(\tilde{V}))$, $y = \pi \circ \phi^{-1}(x)$, $y \in \tilde{P}$. Then $|\sigma - \sigma_D^*| \rightarrow 0$ as $H \rightarrow 0$.

The hypotheses of theorem 2 require only 0-normal hyperbolicity because the splitting we are after is a splitting of the given bundle $N(\tilde{P})$ and not of $T_{\tilde{P}}(\mathbb{R}^n)$. A similar result holds for $\mathcal{L}_D^s$.

5. Numerical implementation

In this section, a specific numerical implementation of the continuation algorithm is described. The algorithm uses the discrete graph transform of section 4 and may be implemented in a computer program following the outline of the present section. In more detail, a practical numerical approach for solving equations (10), (11) and (12) is proposed. The heart of the problem is solving (10), as well as the second equation in (12), for a point $p \in V$. Since this is a global problem, first a method for finding a neighbourhood in V containing p is given, then a method for refining the neighbourhood to a desired tolerance. The numerical conditioning of these problems and numerical error are discussed. Finally, some important smoothing techniques are mentioned. These are useful for controlling the Lipschitz constant of sections and stabilizing a computation in which non-smooth data appears. In section 6 this implementation is used to compute some invariant curves and surfaces.

5.1. Continuation context

Recall that continuation here means the following. Given homotopic diffeomorphisms $F_s : \mathbb{R}^n \rightarrow \mathbb{R}^n$, $s \in [s_0, s_1]$ and an F_{s_0} -invariant submanifold $V \subset \mathbb{R}^n$, successively increment the homotopy parameter s by Δs , computing an approximation to the F_s -invariant manifold at each step. One *continuation step* is the computation of the $F_{s+\Delta s}$ -invariant manifold given an approximation to the F_s -invariant manifold. The input data to a continuation step are equal to the output data of the previous continuation step. These data are as follows.

Continuation step input/output

- (i) A polyhedron P Lipschitz-near a C^r F_s -invariant submanifold $V \subset \mathbb{R}^n$, $r \geq 1$.
- (ii) Approximately normal fibres $N_x(P)$, $x \in \mathcal{C}^0 =$ the vertices of P , and a splitting $N_x(P) = N_x^u(P) \oplus N_x^s(P)$, $x \in \mathcal{C}^0$, which is near the hyperbolic splitting.

Each continuation step has two consecutive stages. The *graph transform algorithm* outputs the data (i). The *linear graph transform algorithm* outputs the data (ii). The linear graph transform algorithm, using $\mathcal{L}_D^u$ and $\mathcal{L}_D^s$ from section 4.3, is not discussed further here. It is less complicated than the graph transform algorithm since it presents no additional nonlinear equations to solve.

5.2. The discrete graph transform algorithm

The graph transform algorithm iterates the graph transform step until the convergence criteria are met. The *graph transform step* takes as input a discrete section σ_D^i of $Z(P)$ and returns as

output a discrete section $\sigma_D^{i+1} = \Gamma_D \circ \sigma_D^i$ of $Z(P)$. Here, $Z(P) = \{(x, v) \in N(P) : |v|_x \leq \epsilon\}$ is from section 4.1 and Γ_D is from section 4.2. The graph transform algorithm starts with the zero section σ_D^0 of $Z(P)$ and for $i \geq 0$ repeats (graph transform step) until the convergence criteria are met. The convergence criteria for the graph transform are the following. The iteration of (graph transform step) is stopped when $|\sigma_D^{i+1} - \sigma_D^i| < \text{error}$ and the contraction factor $|\sigma_D^{j+2} - \sigma_D^{j+1}| / |\sigma_D^{j+1} - \sigma_D^j| < 1$ is approximately constant for all $j < i$ sufficiently large. Similar convergence criteria are discussed in [4].

The graph transform step consists of the following. Recall that Σ_i , defined in section 3.4, is the principal lattice of order $p \geq 1$ of the d -simplex $\mathcal{C}_i$. A discrete section of $Z(P)$ is determined by a discrete set of data points, one in each fibre $Z_x(P)$, $x \in G = \bigcup\{\Sigma_i : i = 1 \dots N\} \subset P$. Thus, for the graph transform step, the *input* is the set of data points $\sigma_D^i(x)$, $x \in G$, and the *output* is the set of data points $\sigma_D^{i+1}(x) = (\Gamma_D \circ \sigma_D^i)(x)$, $x \in G$. The sections have stable and unstable parts, $\sigma_D^i(x) = (x, v^{s,i}(x), v^{u,i}(x))$ and $\sigma_D^{i+1}(x) = (x, v^{s,i+1}(x), v^{u,i+1}(x))$, where $v^{s,i}(x), v^{s,i+1}(x) \in Z_x^s(P) = N_x^s(P) \cap Z_x(P)$ and $v^{u,i}(x), v^{u,i+1}(x) \in Z_x^u(P) = N_x^u(P) \cap Z_x(P)$. Hence, the graph transform step has two independent stages, one for determining the stable part $v^{s,i+1}(x)$, $x \in G$ and another for determining the unstable part $v^{u,i+1}(x)$, $x \in G$.

Now the notation of sections 2 and 3, $\tilde{F} = F_{s+\Delta_s}$ and $F = F_s$, is used. Other notations used below are $\phi : Z(P) \rightarrow U \subset \mathbb{R}^n$, defined in section 2.1, and $\tilde{F}^0 = \phi^{-1} \circ \tilde{F} \circ \phi$, defined in section 2.2.

Graph transform step. Stable part

For $x \in G$:

1. Put $v^s = v^{s,i}$, $v^u = v^{u,i}$ in (10) and (11).
2. Solve (10) for $p \in P$.
 - 2.1 Determine a neighbourhood containing $p \in P$.

$A_j \equiv \bigcup\{\mathcal{C}_k : \mathcal{C}_k \cap \mathcal{C}_j \neq \emptyset\}$ for $j = 1 \dots N$.

Find $j^* \in \{1 \dots N\}$ with $\tilde{F}^0 \circ \sigma_D^i(A_{j^*}) \cap Z_x(P) \neq \emptyset$.

(a) $\mathcal{C}_j^0 \equiv$ vertices of $\mathcal{C}_j$, $j = 1 \dots N$.

(b) $B_j \equiv d$ -simplex with vertices $\phi \circ \tilde{F}^0 \circ \sigma_D^i(\mathcal{C}_j^0)$, $j = 1 \dots N$.

(c) For $j = 1 \dots N$: Test $B_j \cap \phi(Z_x(P)) \neq \emptyset$. If true, return $j = j^*$.
 - 2.2 Locate $p \in A_{j^*}$ to a desired tolerance.

(a) Search for p in each $\mathcal{C}_k \subset A_{j^*}$ using a standard root finding method [16].

(b) If no root found in (a), search $\mathcal{C}_k$ in successively larger regions around A_{j^*} .
3. Evaluate (11) at p to obtain $v^{s,i+1}(x) = w^s(x)$.

In 2.1, a simple geometrical test is used to find A_{j^*} . For each $j = 1 \dots N$, this test determines whether or not a d -simplex intersects an affine $(n-d)$ -plane. This step is typically only necessary for $i = 0$, the same j^* may be used for $i > 0$, since the location of $p \in P$ may not change much as i increases. The approach in 2.1 is justified by the fact that σ_D^i is kept approximately flat over $\mathcal{C}_j$ and $\tilde{F}^0$ is well approximated by its linearization over the set $\sigma_D^i(\mathcal{C}_j)$ as $H \rightarrow 0$.

Graph transform step. Unstable part

For $p \in G$:

1. Put $v^s = v^{s,i}$, $v^u = v^{u,i}$ in (12).
2. Solve (12) for $w = w^u(p) \in Z_p^u(P)$.

Comment: Use a standard root finding method [16] with initial guess $w = 0$. Function evaluations in the root finding method require a call to the following subroutine.

2.1. Given $w \in Z_p^u(P)$, solve the second equation in (12) for $x = x(w) \in P$.

(a) $y \equiv \phi \circ \tilde{F}^0(p, v^{s,i}(p), w)$.

(b) $x \in P$ is the point near y with $y - x$ parallel to $\phi(Z_x(P))$. There are two stages to solving for x , similar to *Stable part* step 2.

3. Put $v^{u,i+1}(p) = w$.

5.3. Practical numerical considerations

Some remarks regarding numerical conditioning [16] are in order. Solving (10) for $p \in P$ is numerically optimally conditioned for $N(P)$ chosen perpendicular to V , as is the evaluation of the second equation of (12). In practice, $N(P)$ is an approximate normal bundle in the following sense. Each d -simplex subspace P_i , $i = 1 \dots N$, of P is a manifold with boundary with tangent bundle $T(P_i)$. Then $\inf\{\angle N_x(P), T_x(P_i) : \text{all } P_i \text{ containing } x, x \in P\} \rightarrow \pi/2$ as $H \rightarrow 0$. In the evaluation of (11) at p , hyperbolicity damps the numerical discretization error and the rounding error, as it does in the first equation of (12). Solving (12) for w^u is a well-conditioned problem. This is because the normal hyperbolicity of V implies that small errors in w^u produce large deviations in the right-hand side of the first equation of (12).

As discussed in section 4.2, it is important to control the Lipschitz constant of discrete sections $\sigma_D(x) = (x, v^s(x), v^u(x))$, $x \in P$. The Lipschitz constant of sections is effectively controlled in practice using two smoothing techniques.

The first smoothing technique is an even redistribution of the grid points G . This replaces P with a nearby polyhedron P' with each $C_i \subset P'$ close to the shape of the standard d -simplex. Redistribution of the grid points is necessary between continuation steps in some cases due to the geometry of the manifold. The distribution of the grid points is not influenced by the dynamics on the manifold.

The second smoothing technique is local fairing [13] of the data $v^s(x) \in N_x^s(P)$ and $v^u(x) \in N_x^u(P)$, $x \in \Sigma_i$, which smooths out $\text{graph}\{\sigma_D\}$. Consider, for example, the attracting case. Here, the data $\sigma_D^i(x) \in Z_x(P)$, $x \in \Sigma_i$, are tested for large deviations. If an undesirable data point $\sigma_D^i(x^*)$ is detected, it is replaced by the average of $\sigma_D^i(x)$, $x \neq x^*$, $x \in \Sigma_i$. To be precise, the average $y \in \mathbb{R}^n$ of $\phi \circ \sigma_D^i(x) \in \mathbb{R}^n$, $x \neq x^*$, $x \in \Sigma_i$, is obtained. Then, y is projected onto the affine k_1 -plane $\phi \circ Z_x^*(P)$ to obtain $z \in \phi \circ Z_x^*(P) \subset \mathbb{R}^n$. The data point $\sigma_D^i(x^*)$ is replaced by $\phi^{-1}(z)$.

Between continuation steps, it is also important to use local averaging of the fibres of $N(P)$, to make $N_x(P)$, $x \in C_i$, more nearly parallel. For each $x \in C^0$, $N_x(P)$ is replaced by the average of the $N_y(P)$ for $y \in C^0$ near x . This is sometimes necessary because, in practice, small bumps in P can introduce degeneracies in its approximate normal bundle $N(P)$.

Remarks

- (i) Recall from section 3 that σ_D is constructed from data points $\sigma_D(x)$, $x \in G$, simplex-wise, as follows. For $x \in C_i$,

$$\sigma_D(x) = (x, \xi_i^s(x) \circ \eta_i^s(x), \xi_i^u(x) \circ \eta_i^u(x)).$$

From sections 3.2 and 3.3, $\xi_i^s(x) : \mathbb{R}^{k_1} \rightarrow N_x^s(P)$, $\xi_i^s(x)(\eta) = E_i^s(x)\eta$, where the columns of $E_i^s(x)$ give an orthonormal basis for $N_x^s(P)$, continuous for $x \in C_i$. The formulation of $\xi_i^u(x)$ is analogous. Now, the data points $\sigma_D(x_{i,j}) = (x_{i,j}, v_{i,j}^s, v_{i,j}^u) \in N_{C_i}(P)$, where

$x_{i,j} \in \mathcal{C}_i$, $j = 1 \dots m$, are the points of the principal lattice Σ_i . The polynomial $\eta_i^s(x)$ is constructed from $v_{i,j}^s$, $j = 1 \dots m$. There is a unique Lagrange polynomial $\eta_i^s : \mathcal{C}_i \rightarrow \mathbb{R}^{k_1}$ of total degree p fitting the data

$$\eta_i^s(x_{i,j}) = \xi_i(x_{i,j})^{-1}(v_{i,j}^s), \quad j = 1 \dots m.$$

The computation of these polynomials is standard [7, 9, 16]. The construction of $\eta_i^u(x)$ is analogous.

- (ii) The size of the continuation step Δs may be guessed by estimating the size of $\theta = \|F - \tilde{F}\|_{C^1}$ in section 2.2. This is done by noting the restrictions on the size of θ at each stage of the proof of the invariant manifold theorem. Since the proof is completely general, in many cases the estimates obtained in this way are too stringent in practice. Nevertheless, such estimates are an important qualitative guide to guessing the continuation step. For simplicity, consider the case of V attracting. The starting condition which sets the stage for the rest of the proof is

$$\tilde{F}^0(Z(\epsilon)) \subset Z(\epsilon) \quad \text{for some } \epsilon > 0. \quad (14)$$

Here, $Z(\epsilon)$ is from section 2.1 and $\tilde{F}^0$ is from section 2.2. This condition is a consequence of the normal contraction towards V under F and the smallness of θ . Suppose the normal contraction towards V under F is measured by $A \in (0, 1)$, where

$$F^0(Z(\epsilon)) \subset Z(A \cdot \epsilon) \quad \text{for some } \epsilon > 0. \quad (15)$$

Then (14) is implied by $|\tilde{F} - F|_{C^0} < (1 - A)\epsilon$. Here, A is an approximation to the linear normal contraction and the smallness of ϵ depends on nonlinear terms. If (15) is known for some A, ϵ , perhaps from numerical tests, then $\theta < (1 - A)\epsilon$ is one estimate of θ .

It may be desirable to have an estimate of θ which does not depend on ϵ , but only on the initial problem parameters. This is accomplished by linearizing F at V . Suppose the linear normal contraction towards V under F is $a \in (0, 1)$ where $\|DF|N_y^s(V)\| < a$, $y \in V$. Then

$$F^0(Z(\epsilon)) \subset Z((a + c_2\epsilon/2) \cdot \epsilon),$$

where $c_2 = \sup\{\|D^2F(y)\| : y \in \phi(Z(\epsilon))\}$. Thus, (14) is implied by $|\tilde{F} - F|_{C^0} < (1 - a - c_2\epsilon/2)\epsilon$, where $\epsilon < 2(1 - a)/c_2$. Choosing $\epsilon = (1 - a)/c_2$, we obtain that $|\tilde{F} - F|_{C^0} < (1 - a)^2/(2c_2)$ implies (14). In conclusion, one estimate of θ is $\theta < (1 - a)^2/(2c_2)$.

Restrictions on the size of θ at later stages of the proof are similar. For example, one upper bound on θ is directly proportional to $1 - b$, where $b = c\mu^N \in (0, 1)$ measures the strength of the normal hyperbolicity. Here, c and μ are from (2) and N is from the end of section 2.2. The upper bounds on θ are also inversely proportional to powers of $c_1 = \sup\{\|DF(y)\|, \|(DF)^{-1}(y)\| : y \in \phi(Z(\epsilon))\}$. The estimates derived at these later stages, while theoretically valid, appear to be too stringent in practice.

6. Computations

This section contains some examples of computations performed with the algorithm described in section 5. There are a variety of test cases in section 6.1. Section 6.2 contains an application to computing the ‘slow-transient’ surface of an enzyme reaction model. This application requires a modification to the algorithm of section 5 which allows the computation of just a part of an invariant manifold. All examples in section 6.1 compute the whole invariant manifold. In section 6.2, a piece of an invariant manifold is computed. This is a necessary

adaptation in cases where the invariant manifold is so large that its data cannot be held in computer memory.

The smoothing techniques of section 5.3 are not necessary in all examples. Redistribution of the grid points and local averaging of the normal fibres are used when necessary between continuation steps in all cases. However, local fairing of the data is not used in the examples of section 6.1.

The convergence theory of Γ_D , $\mathcal{L}_D^u$ and $\mathcal{L}_D^s$ in section 4 is borne out by the experiments of the present section. In particular, the contraction factors of these operators agree in practice with the theoretical values. Although the convergence theory of section 4 incorporates discretization effects, in actual computations there is an issue of numerical resolution. That is, computer memory and processor limits impose constraints on the size of the grid of the discrete manifold. In the case of vector fields, there is also a limit on the numerical resolution of the integrated vector field. The numerical resolution is measured by the discretization error tolerance $\mathcal{E} > 0$, the sum of the discrete manifold approximation error and (if applicable) the vector field integration error. In practice, the hyperbolicity damps the error from these sources.

From the numerical experiments, the following general statement may be made regarding the observations of the empirical behaviour of the algorithm. To simplify this, suppose the perturbed manifold $\tilde{V}$ is attracting. Given appropriate initial data, the algorithm converges to the fixed point of Γ_D , to arbitrary accuracy, provided the normal contraction factor $\alpha < 1$ does not approach 1. This convergence is fairly insensitive to the size of the grid of the discrete manifold. Theoretical estimates suggest convergence in practice provided α does not approach $1 - \mathcal{E}/\epsilon$, $\mathcal{E}/\epsilon < 1$. Here, ϵ is the diameter of the tubular neighbourhood, which measures the size of the contracting neighbourhood of V . Numerically, the factor $1/\epsilon$ scales the discretization error $\mathcal{E}$ to the region of interest. That is, $\mathcal{E}/\epsilon$ is the relative numerical error in computing $\tilde{V}$. In the extreme case $\alpha \rightarrow 1$, convergence of the algorithm requires that $\mathcal{E}$ is small compared with $1 - \alpha$.

As indicated in section 5.2, the innermost loop of the procedure, where the main numerical work is done, consists of a standard local root finding method. This is in *stable part* step 2.2(a) and *unstable part* step 2.1(b). In the present section, Newton's method is used here, with a tolerance on function values of $1.0e-6$. The convergence of the graph transform algorithm is relatively insensitive to the tolerance of the root finding method. This is because the error in the root is tangent to the manifold. Thus, the error is mostly annihilated upon evaluation of (11) or the first equation in (12).

6.1. Examples

This section contains some of the numerical experiments performed to test the continuation algorithm. The present test suite of dynamical systems was chosen to validate the algorithm in a variety of cases. Thus, the examples have either been computed by alternate means elsewhere or theoretical analysis indicates the expected qualitative appearance of the manifold.

To demonstrate that the algorithm converges regardless of the dynamics on the manifold, examples were chosen with a variety of dynamics. The first three examples are computations of invariant tori. In the first torus there is a dense intersection of the stable and unstable manifolds of a saddle point. In the second torus, there are two saddle points, a source and a sink—almost all the points on the torus are in the basin of attraction of the sink. The third is a quasiperiodic torus of a vector field. The last two examples are saddle type invariant curves. In the first of these there is a range of dynamics as the continuation parameter varies: initially quasiperiodic with phase locking, then a saddle-node bifurcation. The second is a saddle type periodic orbit

Figure 4. Invariant tori, 5000 simplices. (Left) The fattened Thom map, $\epsilon = 0.469$, $p = 3$. (Right) The fattened sink map, $\epsilon = 0.353$, $p = 1$.

of a vector field. No special difficulty was observed in the continuation due to these different dynamical scenarios.

Examples 1, 2 and 3 are attracting and do not test the linear graph transform. Examples 1 and 2 test the graph transform in the diffeomorphism case and example 3 tests the vector field case. The saddle case graph transform and the linear graph transforms which compute the stable and unstable bundles are tested in examples 4 and 5. Example 4 tests the diffeomorphism case and example 5 tests the vector field case. In examples 1, 2 and 4 the given systems are split over an analytically known initial manifold. On the other hand, in examples 3 and 5 there is no analytically known initial manifold for the continuation. In examples 3 and 5, initial data are obtained by simulation and are quite rough. Here, Γ_D converged with this rough initial data, smoothing it out. Examples 1 and 4 illustrate that Γ_D remains contractive in practice when the discrete manifold is of order $p > 1$.

In addition, the most informative tests of an algorithm's empirical behaviour may be at the limits of performance. Thus, most of the test problems explore the performance of the algorithm with one of the algorithm's parameters near a theoretical limit. Usually this has to do with a lack of hyperbolicity at one or more points. This may or may not be accompanied by an observable loss of smoothness of the manifold.

Example 1. The fattened Thom map [4]

$$\begin{aligned} x_{i+1} &= 0.1x_i + \epsilon \sin z_i, \\ y_{i+1} &= z_i + y_i + \epsilon x_i, & (x_i, y_i, z_i) &\in \mathbb{R} \times \mathbf{S}^1 \times \mathbf{S}^1, \\ z_{i+1} &= 2z_i + y_i + \epsilon x_i. \end{aligned} \quad (16)$$

This diffeomorphism is a dynamical system defined on $\mathbb{R} \times \{\text{the standard 2-torus}\}$ or $\mathbb{R} \times \mathbf{S}^1 \times \mathbf{S}^1$. The usual coordinates are used to identify $\mathbf{S}^1$ with $\mathbb{R}/2\pi\mathbb{Z}$.

For $\epsilon = 0$, $\{x = 0\}$ is an attracting torus and there is one hyperbolic fixed point $(0, 0, 0)$. The stable and unstable manifolds of this saddle point, restricted to the torus $\{x = 0\}$, have dense intersection in the torus. Thus, the system exhibits hyperbolic mixing. In the present example, the attracting torus is computed for $\epsilon \geq 0$. For all $\epsilon > 0$, the point $(0, 0, 0)$ remains a fixed point. The eigenvalues go from about 0.100, 0.382, 2.62 at $\epsilon = 0$ to approximately 0.192, 0.192, 2.72 at about $\epsilon = 0.469$, at which point the deformed torus fails to be normally hyperbolic.

In the present example, the discrete manifold was of order $p = 3$ and the polyhedron P contained $50 \times 50 \times 2 = 5000$ d -simplices. The initial data consisted of the standard 2-torus $V = \{x = 0, 0 \leq y < 2\pi, 0 \leq z < 2\pi\}$ together with the normal fibres $N_p(V) = \text{span}\{(1, 0, 0)\}$ at grid points $p = (0, y, z) \in G$.

Starting from $\epsilon = 0.0$, the attracting torus was continued to $\epsilon = 0.469$ using ϵ steps of 0.1 and a tolerance of $\text{error} = 1.0e - 12$. The last ϵ step was 0.069. See figure 4. The average

contraction during the continuation was 0.126, close to 0.1, which is what we expect from (16). In a separate series of computations, the algorithm was started with initial data $\{x = 0\}$ three times. In these three runs, the parameter ϵ was set to 0.3, 0.4 and 0.469. The number of graph transform iterates required for convergence to tolerance $error = 1.0e - 12$ was 13, 17 and 19.

This map was introduced in [4], where the same surface was computed using a different algorithm based on the graph transform. For a computation of this surface using a discrete manifold of order $p = 1$ together with adaptive non-uniform refinement of the grid, see [2].

Example 2. The fattened sink map

$$\begin{aligned} x_{i+1} &= 0.25x_i + \epsilon \sin z_i, \\ y_{i+1} &= y_i + 0.5 \sin y_i + \epsilon x_i, \\ z_{i+1} &= z_i - 0.5 \cos y_i \sin z_i + \epsilon x_i. \end{aligned} \quad (x_i, y_i, z_i) \in \mathbb{R} \times \mathbf{S}^1 \times \mathbf{S}^1, \quad (17)$$

This map is a diffeomorphism for $\epsilon \in [0, 1/\sqrt{8}) \approx [0, 0.353)$. It is a dynamical system defined on $\mathbb{R} \times \mathbf{S}^1 \times \mathbf{S}^1$. The usual coordinates are used to identify $\mathbf{S}^1$ with $\mathbb{R}/2\pi\mathbb{Z}$.

For $\epsilon = 0$, there are four hyperbolic fixed points, $(0, \pi, \pi)$, $(0, 0, \pi)$, $(0, 0, 0)$ and $(0, \pi, 0)$. In the present example, the unstable manifold of $(0, 0, \pi)$ is computed for $\epsilon \geq 0$. At $\epsilon = 0$, the unstable manifold is $\{x = 0, 0 \leq y < \pi, \pi < y < 2\pi, 0 < z < 2\pi\}$. This is the torus $\{x = 0\}$ with the two circles $\{x = 0, z = 0\}$ and $\{x = 0, y = \pi\}$ taken out. The closure of the unstable manifold, $\{x = 0\}$, is a normally hyperbolic (attracting) torus. Almost all the points in the unstable manifold are in the basin of attraction of the sink $(0, \pi, \pi)$.

For all $\epsilon > 0$, the four points mentioned above remain fixed points. The eigenvalues at $(0, 0, \pi)$, which determine the existence of the unstable manifold, go from 0.25, 1.5, 1.5 at $\epsilon = 0$ to 0.75, 1.0, 1.5 at about $\epsilon = 0.612$. The eigenvalues of the sink $(0, \pi, \pi)$ go from 0.25, 0.5, 0.5 at $\epsilon = 0$ to approximately 0.375, 0.375, 0.5 at about $\epsilon = 0.125$, at which point the compact attracting manifold containing the unstable manifold fails to be normally hyperbolic. Thus, for $\epsilon > 0.125$, the unstable manifold is not smooth at the point $(0, \pi, \pi)$. A spiral develops here due to the existence of a pair of complex conjugate eigenvalues whose eigenspace is a plane which intersects the torus transversally.

In the present example, the discrete manifold was of order $p = 1$ and the polyhedron P contained $50 \times 50 \times 2 = 5000$ d -simplices. The initial data consisted of the standard 2-torus $V = \{x = 0, 0 \leq y < 2\pi, 0 \leq z < 2\pi\}$ together with the normal fibres $N_p(V) = \text{span}\{(1, 0, 0)\}$ at grid points $p = (0, y, z) \in G$.

Starting from $\epsilon = 0.0$, the unstable manifold was continued to $\epsilon = 0.353$ using ϵ steps of 0.1 and a tolerance of $error = 1.0e - 12$. The last ϵ step was 0.053. See figure 4. The average contraction through the first two continuation steps was 0.265, close to 0.25 as expected from (17). In separate computations, the algorithm was started with initial data $\{x = 0\}$ three times. In these three runs, the parameter ϵ was set to 0.2, 0.3 and 0.353. The number of graph transform iterates required for convergence to tolerance $error = 1.0e - 12$ was 19, 22 and 24.

Example 3. The Lorenz-84 system [33]

$$\begin{aligned} \dot{x} &= -y^2 - z^2 - 0.25x + 0.25F, \\ \dot{y} &= xy - 4xy - y + G, \\ \dot{z} &= 4xy + xz - z. \end{aligned} \quad (x, y, z) \in \mathbb{R}^3, \quad (18)$$

This vector field is a low-order model of atmospheric circulation on the globe. It is proposed as a simple model possessing the qualitative features of Hadley circulation in [33]. The system is dissipative—the total energy $(x^2 + y^2 + z^2)/2$ decreases if it exceeds a certain value.

Figure 5. Lorenz-84 invariant tori; $p = 1$, 32768 simplices; (left) $F = 1.84$, (right) $F = 1.755$.

Certain curves separate the (F, G) parameter plane into two regions where (18) possesses either one or three fixed points. On one of these curves there is a codimension-2 bifurcation point $(F, G) \approx (1.6840, 1.6829)$, [47]. A system in the vicinity of such a point may give rise to a variety of dynamical behaviour, for example, quasiperiodic solutions and homoclinic bifurcations [30].

We focus on the region of parameter space for which (18) possesses a repelling quasiperiodic torus. In this region there are three fixed points—two of saddle type near the torus and one sink. The torus encircles an attracting periodic orbit. On one boundary of this region of parameter space the attracting periodic orbit becomes repelling and the torus disappears. On the other boundary the torus becomes a heteroclinic connection between the saddle points. Beyond the immediate vicinity of the codimension-2 point the dynamics becomes even more varied. For example, for $(F, G) \approx (2.5, 1.4)$, the torus becomes a chaotic repeller [47].

In the present example, the repelling invariant torus was continued along a segment with G fixed, in both directions starting from $(F, G) = (1.8, 1.65)$. The initial data were obtained by numerical simulation. See remarks (i) and (ii). The discrete manifold was of order $p = 1$ and the polyhedron P contained $128 \times 128 \times 2 = 32768$ d -simplices.

The invariant torus was continued for $F \in (1.755, 1.84)$ using steps of $\Delta F = 0.01$ and a tolerance of $\text{error} = 1.0e - 4$. See figure 5. For decreasing F the torus approaches a more sphere-like surface and the inner radius gets smaller. This is expected since the parameters approach the Hopf saddle-node bifurcation point [30].

For F outside the interval $(1.755, 1.84)$ computational instability was observed, although simulation suggests that the torus continues to exist. This instability and the high tolerance used in this example both highlight the delicacy of this computation. Two main factors make this computation challenging. First, the initial data are not analytically known. It must be roughly estimated by numerical simulation. Second, the attractivity of the torus (for reverse time) appears to be quite weak. This makes sense in light of the dynamics described above. The nearness of the codimension-2 point and widely varied dynamics imply that small perturbations in the (F, G) parameter plane may destroy the torus.

Example 4. The fattened Arnold family [4]

$$\begin{aligned} x_{i+1} &= x_i + 0.1 + \epsilon(y_i + z_i/2 + \sin x_i), \\ y_{i+1} &= 0.3(y_i + \sin x_i), & (x_i, y_i, z_i) &\in \mathbf{S}^1 \times \mathbb{R}^2, \\ z_{i+1} &= 2.4(y_i + z_i + \sin x_i). \end{aligned} \tag{19}$$

This diffeomorphism is a dynamical system defined on $\mathbf{S}^1 \times \mathbb{R}^2$. The usual coordinates are used to identify $\mathbf{S}^1$ with $\mathbb{R}/2\pi\mathbb{Z}$. A similar diffeomorphism is studied in [5].

For $\epsilon = 0$, there is an invariant closed curve of saddle type which is the graph of an explicit function over $\mathbf{S}^1$, [4]. The closed curve is filled with dense quasiperiodic orbits. For small $\epsilon > 0$, the invariant closed curve persists. The closed curve is filled with either quasiperiodic orbits or periodic orbits, depending on the value of ϵ . The regions of periodic behaviour, or phase locking, are described in terms of Arnold tongues [30]. At about $\epsilon = 0.49$, two saddle points appear on the closed curve in a saddle-node bifurcation. With respect to (19) restricted to the closed curve, the fixed points are a sink and a source. At about $\epsilon = 0.776$, the tangential contraction at the sink equals the normal contraction, and the curve loses its normal hyperbolicity.

In the present example, the discrete manifold was of order $p = 3$ and the polyhedron P contained 50 d -simplices. The initial data consisted of the invariant curve together with a hyperbolic splitting for $\epsilon = 0$, for which there were explicit formulae.

Starting from $\epsilon = 0.0$, the closed curve was continued to $\epsilon = 0.776$ using ϵ steps of 0.2. The last ϵ step was 0.176. The tolerance for both the graph transform and the linear graph transform was $error = 1.0e - 12$. The last ϵ value for which the hyperbolic splitting was computed was 0.6. The angle between the stable and unstable parts of the hyperbolic splitting appears to $\rightarrow 0$ for increasing ϵ . The average contraction of the stable part of the graph transform through the continuation was 0.331, close to 0.3, as expected from (19). The average contraction of the unstable part of the graph transform through the continuation was 0.425, close to $1.0/2.4 \approx 0.417$, as expected from (19). The average contraction of the linear graph transform for the unstable bundle, through the continuation, was 0.128. The average contraction of the linear graph transform for the stable bundle, through the continuation, was 0.129. Both these contraction rates are close to $0.3/2.4 = 0.125$, as expected from (19). See remark (iii).

This diffeomorphism was introduced in [4], where this closed curve and its hyperbolic splitting were also computed.

Example 5. The Lorenz system [32]

$$\begin{aligned}\dot{x} &= \sigma(y - x), \\ \dot{y} &= rx - y - xz, \\ \dot{z} &= xy - bz.\end{aligned}\quad (x, y, z) \in \mathbb{R}^3, \quad (20)$$

This vector field is a three-mode approximation to the motion of a layer of fluid heated from below. Originating in the context of numerical weather prediction, it was discovered in a search for a simple system possessing bounded aperiodic solutions [32]. The system is dissipative—the Liapunov function $rx^2 + \sigma y^2 + \sigma(z - 2r)^2$ decreases if it exceeds a certain value. Thus, the system has a unique global attractor for all parameter values $\sigma, b, r > 0$.

We focus on the well-studied case $\sigma = 10.0$ and $b = 8/3$. For $r \in (0, 1)$, the equilibrium $(0, 0, 0)$ is a global attractor. At $r = 1$, this equilibrium undergoes a pitchfork bifurcation. Thus, $(0, 0, 0)$ becomes a saddle point and two attracting equilibria appear, $x = y = \pm[b(r - 1)]^{1/2}$, $z = r - 1$. The basins of attraction of these equilibria are separated by the codimension-1 stable manifold of $(0, 0, 0)$. At about $r = 13.962$, the unstable manifold of $(0, 0, 0)$ forms a pair of homoclinic loops. These loops indicate a non-transverse intersection of the stable and unstable manifolds of $(0, 0, 0)$. For $r > 13.962$, a strange invariant set appears due to a homoclinic explosion. This invariant set consists of a countable infinity of periodic orbits, an uncountable infinity of aperiodic orbits and an uncountable infinity of orbits terminating at $(0, 0, 0)$. At about $r = 24.74$, the two attracting equilibria undergo subcritical Hopf bifurcations. For $r > 24.74$ all three equilibria are saddles and the Lorenz attractor appears. The solutions on both the strange invariant set and the Lorenz attractor

exhibit aperiodic behaviour and sensitive dependence on initial conditions. Most solutions go back and forth between neighbourhoods of two equilibria in an apparently unpredictable fashion [38, 49].

The two periodic orbits arising from the Hopf bifurcations, for $r \in (13.962, 24.74)$, tend towards the homoclinic loops for decreasing r . In the present example, one of these saddle type periodic orbits is continued for decreasing r starting at $r = 20.0$. This orbit is contained in the strange invariant set, thus also in the global attractor. The initial data were obtained by numerical simulation. In particular, the initial hyperbolic splitting was obtained as follows. The variational equation was integrated to approximate the derivative of a Poincaré map. The approximate derivative was then numerically solved for the eigenvectors. The discrete manifold was of order $p = 1$ and the polyhedron P contained 200 d -simplices.

The periodic orbit was continued to $r = 16.5$ using r steps of 0.1. The tolerance for both the graph transform and the linear graph transform was $error = 1.0e - 12$. Figure 1 shows the final curve near a homoclinic loop. In this example it was necessary to redistribute the grid points near the sharp bend in the curve to prevent their spread during continuation. To continue the orbit further, a discrete manifold with either more d -simplices or higher order is required. This is necessary to resolve the sharp bend in the orbit near $(0, 0, 0)$.

Remarks

- (i) Numerical continuation allows the approximation of an F_{s_1} -invariant manifold V_{s_1} by starting from an analytically known F_{s_0} -invariant V_{s_0} . This works in principle as long as there are normally hyperbolic F_s -invariant V_s for $s \in [s_0, s_1]$. Even if no V_{s_0} is known analytically, it may still be possible to obtain starting data. It is enough to be able to estimate, numerically or otherwise, a P sufficiently close to V_{s_0} . In the attracting case, the only other starting data is $N(P)$. There are standard methods for computing orthonormal bases for $N_x(P)$ from any d -simplex of P containing x , $x \in \mathcal{C}^0$, using the singular value decomposition [40].
- (ii) In example 3, the starting torus was obtained in the following way. First, a large set of points S in $\mathbb{R}^3$ near the torus was obtained by following trajectories and discarding the initial parts of the trajectories. To obtain a uniform grid of points, coordinates were introduced as follows. A circle lying inside the torus-shape made by S was estimated. The coordinates of a point in S consisted of an angle θ on this circle and polar coordinates (ϕ, r) in the plane orthogonal to the circle at θ . Fix $N \geq 1$ (in our case $N = 128$), $\Delta\theta > 0$ small and $\Delta\phi > 0$ small. For each $i, j = 0, \dots, N - 1$, a point in S was found whose (θ, ϕ, r) -coordinates satisfy $\theta \in [2\pi i/N, 2\pi i/N + \Delta\theta]$ and $\phi \in [2\pi j/N, 2\pi j/N + \Delta\phi]$. Denote the r -coordinate of such a point by $r_{i,j}$. The set $(\theta, \phi, r) = (2\pi i/N, 2\pi j/N, r_{i,j})$, $i, j = 0, \dots, N - 1$, is then an equally spaced grid of points close to the invariant torus.
- (iii) Here the theoretical contraction rate for the discrete linear graph transform is discussed. The (un)stable bundle is a subbundle of the *a priori* known approximate normal bundle $N(\tilde{P})$. Thus, the contraction rate depends on the stable contraction and unstable expansion factors associated with the hyperbolic splitting—but not on the tangent contraction/expansion factors. Suppose $\alpha \in (0, 1)$ is the stable contraction factor and $\gamma \in (1, \infty)$ is the unstable expansion factor. Then the contraction rate of $\mathcal{L}_D^u$, which determines $N^u(\tilde{P}) \subset N(\tilde{P})$, is α/γ . The contraction rate of $\mathcal{L}_D^s$, which determines $N^s(\tilde{P}) \subset N(\tilde{P})$, is also α/γ . For comparison, this rate is contrasted with the contraction rate of the linear graph transform for determining the $D\tilde{F}$ -invariant unstable bundle in $T_{\tilde{P}}(\mathbb{R}^n)$. Suppose $\beta \in (1, \infty)$ is the bound on the tangent expansion factor.

The contraction rate for determining the *invariant* unstable bundle in $T_{\tilde{p}}(\mathbb{R}^n)$ is α/β . By the definition of r -normal hyperbolicity, $\beta^r < \gamma$ and so $\alpha/\gamma < \alpha/\beta$.

- (iv) There are many alternative approaches to computing one-dimensional invariant manifolds such as periodic orbits, invariant closed curves, heteroclinic and homoclinic orbits. A review of these is outside the scope of the present paper, which presents an algorithm applicable to the case of general higher dimensional manifolds. In light of examples 4 and 5, the following remark is made. There are classical alternative approaches to computing saddle type periodic orbits of vector fields. These often involve finding a hyperplane normal to the orbit and using Newton's method to find the fixed point of the computed Poincaré map. There are also algorithms for computing saddle type invariant closed curves of a map which have been around for a while [25].
- (v) As mentioned in remark (iii) of section 1, two diffeomorphic invariant manifolds $\tilde{V}_1$ and $\tilde{V}_2$ may coexist near each other. On the other hand, there is a neighbourhood U of $\tilde{V}_1$ within which there is no other invariant manifold Lipschitz-close to $\tilde{V}_1$. The size of the neighbourhood U generally depends on the strength of the hyperbolicity. The possibility of coexistence means that careful computation requires the continuation step size to be small, see remark (ii) of section 5. The computations of the present paper were repeated with several different continuation parameter step sizes. The specific computations included here have moderate step sizes.

6.2. An application

This section deals with a problem of chemical kinetics. The 'slow-transient' surface of an enzyme reaction is computed for a variety of parameter values. This invariant surface, in the phase space of chemical species concentration variables, is useful in chemical kinetics for model reduction. Specific knowledge of the surface can also help to quantify the evolution of the reaction. After a short time interval, the n -tuple of chemical species concentrations is restricted to the surface, at least for experimentally measurable tolerances. The dynamics of the reaction after this short time interval is described by the dynamics on the surface. In principle, once this surface is known, the system may be reduced to a 2D system on the surface. In chemical kinetics, the steady state and equilibrium approximations, as well as variations on these, have been used to approximate the slow-transient surface [15]. These approximations are typically valid in limiting cases.

As mentioned in the introduction of section 6, a modification to the algorithm of section 5 is required for the present computation. This is because what is computed is a manifold with boundary S , only part of an attracting surface and not overflowing invariant. For a diffeomorphism F , a compact manifold with boundary V is overflowing invariant under F if $V \subset F(V^0)$, where $V^0 = V \setminus \partial V$ is the interior of V . For a vector field, V is overflowing invariant if the vector field points outwards at every point of the boundary of V . For attracting overflowing manifolds, the graph transform works with no modification [14]. In the present example, the graph transform is not directly applicable to computing S . This is because the pre-image of points in a small tubular neighbourhood of S (but close to the boundary) may not lie in a tubular neighbourhood of S . This issue is addressed by using local extrapolation of S at its boundary after each graph transform step. This means the following. In the present case, the order of approximation is $p = 1$. Thus, the output data of a graph transform step determine σ_D^i where the $\text{graph}\{\sigma_D^i\} = P$ is a polyhedral manifold with boundary. The d -simplices of P whose points are on the boundary of P are flatly extended to form a slightly larger polyhedron $P' \supset P$. This P' is used as the input to the next graph transform step. There

Figure 6. Enzyme reaction surfaces: (left) $k_p = 0.1, k_1 = 10^3$; (middle) $k_p = 0.1, k_1 = 1.0$; (right) $k_p = 1.0, k_1 = 1.0$.

is a history of different approaches to computing slow-transient surfaces in chemical kinetics and non-equilibrium statistical physics, see, for example, [17–19, 42].

In the enzyme reaction model

$$\begin{aligned}\dot{s} &= -k_1(e_0 - c - q)s + k_{-1}c, \\ \dot{c} &= k_1(e_0 - c - q)s - (k_{-1} + k_2)c + k_{-2}q, \\ \dot{q} &= k_2c - (k_{-2} + k_p)q,\end{aligned}\quad (s, c, q) \in \mathbb{R}^3, \quad (21)$$

the variables s, c and q are the concentrations of different chemical species undergoing chemical reaction [42]. Here, $k_1, k_{-1}, k_2, k_{-2}, k_p > 0$ are the rate constants and $e_0 > 0$ is the concentration of the enzyme, taken to be constant. The attracting equilibrium is 0 in the physical region $\{0 \leq s < \infty, c + q \leq e_0, 0 \leq c, q\} \subset \mathbb{R}^3$. In figure 6, the part of the slow-transient surface in the physical region restricted to $\{0 \leq s \leq 2\}$ is computed for three parameter choices. In every case, $e_0 = 1.0, k_{-1} = 1.0, k_2 = 1.0$ and $k_{-2} = 1.0$.

In the present example, the dynamics are described by a nested hierarchy of attracting invariant manifolds in 3D. This is an equilibrium point contained in a curve contained in a surface, the slow-transient surface, which separates the physical region of phase space. The rate of attraction towards the surface is faster than towards the curve in the surface. The rate of attraction towards the curve in the surface is faster than towards the point in the curve. The part of the slow-transient surface in the physical region restricted to $\{0 \leq s \leq 2\}$ is a manifold with boundary S .

The three approximations of S in figure 6 contain 4096 2-simplices. The initial data for the left and the middle surfaces are the surface $c = 0$. One continuation step is used: the left surface is obtained after five graph transform steps, the middle after 47 graph transform steps. This disparity in the rate of convergence illustrates a property of Γ_D discussed in section 4.2. Namely, the contraction factor of Γ_D is directly proportional to the factor of normal contraction towards V under F . That is, the normal contraction is stronger for the left surface than for the middle surface. The right surface is obtained from the middle surface by continuation in k_p where $\Delta k_p = 0.1$. The middle surface is computed by alternate means in [42].

Acknowledgments

This work is supported by the Netherlands Organisation for Scientific Research (NWO), Project No 613-02-201. The output curve and surface data have been displayed using Geomview [39].

References

- [1] Boothby W M 1975 *An Introduction to Differentiable Manifolds and Riemannian Geometry* (New York: Academic)
- [2] Broer H W, Hagen A and Vegter G 2003 Multiple purpose algorithms for invariant manifolds *Dyn. Contin. Discrete Impl. Syst. B* **10** 331–44
- [3] Broer H W, Hagen A and Vegter G 2003 Numerical approximation of normally hyperbolic invariant manifolds *Proc. 4th AIMS Meeting 2002 at Wilmington, DCDS 2003 supplement volume* ed S Hu *et al* (Springfield, MO: AIMS) pp 133–40
- [4] Broer H W, Osinga H M and Vegter G 1997 Algorithms for computing normally hyperbolic invariant manifolds *Z. Angew. Math. Phys.* **48** 480–524
- [5] Broer H W, Simó C and Tatjer J C 1998 Towards global models near homoclinic tangencies of dissipative diffeomorphisms *Nonlinearity* **11** 667–770
- [6] Cairns S 1961 A simple triangulation method for smooth manifolds *Bull. Am. Math. Soc.* **67** 389–90
- [7] Carey G and Oden J 1984 *Finite Elements* vol 3 (Englewood Cliffs, NJ: Prentice-Hall)
- [8] Castellà E and Jorba À 2000 On the vertical families of two-dimensional tori near the triangular points of the bicircular problem *Celest. Mech. Dyn. Astron.* **76** 35–54
- [9] Ciarlet P G and Raviart P 1972 General Lagrange and Hermite interpolation in $\mathbb{R}^n$ with applications to finite element methods *Arch. Ration. Mech. Anal.* **46** 177–99
- [10] Dellnitz M, Froyland G and Junge O 2001 The algorithms behind GAIO-set oriented numerical methods for dynamical systems *Ergodic Theory, Analysis and Efficient Simulation of Dynamical Systems* ed B Fiedler (Berlin: Springer)
- [11] Dieci L and Bader G 1994 Solution of the systems associated with invariant tori approximation: II. Multigrid methods *SIAM J. Sci. Comput.* **15** 1375–400
- [12] Dieci L and Lorenz J 1995 Computation of invariant tori by the method of characteristics *SIAM J. Numer. Anal.* **32** 1436–74
- [13] Farin G 1997 *Curves and Surfaces for Computer-Aided Geometric Design: A Practical Guide* (New York: Academic)
- [14] Fenichel N 1971 Persistence and smoothness of invariant manifolds for flows *Indiana Univ. Math. J.* **21** 193–226
- [15] Fraser S J 1988 The steady state and equilibrium approximations: a geometrical picture *J. Chem. Phys.* **88** 4732–8
- [16] Golub G and Ortega J M 1993 *Scientific Computing: An Introduction With Parallel Computing* (San Diego, CA: Academic)
- [17] Gorban A N and Karlin I V 2005 *Invariant Manifolds for Physical and Chemical Kinetics (Lecture Notes in Physics* vol 660) (Berlin: Springer)
- [18] Gorban A N, Karlin I V and Zinovyev A Yu 2004 Constructive methods of invariant manifolds for kinetic problems *Phys. Rep.* **396** 197–403
- [19] Gorban A N, Karlin I V and Zinovyev A Yu 2004 Invariant grids for reaction kinetics *Physica A* **333** 106–54
- [20] Hagen A 1996 Hyperbolic structures of time discretizations and the dependence on the time step *PhD Thesis* University of Minnesota
- [21] Hagen A 2004 Hyperbolic trajectories of time discretizations *Nonlinear Anal.* **59** 121–32
- [22] Haro A and De La Llave R 2005 A parametrization method for the computation of invariant tori and their whiskers in quasi-periodic maps: numerical implementation and examples *UB-UPC Preprint 13*
- [23] Haro A and De La Llave R 2004 A parametrization method for the computation of invariant tori and their whiskers in quasi periodic maps: II. Numerical algorithms *UB-UPC Preprint 11*
- [24] Haro A and De La Llave R 2004 A parametrization method for the computation of invariant tori and their whiskers in quasi periodic maps: I. Rigorous results *UB-UPC Preprint 10*
- [25] Kevrekidis I, Aris R, Schmidt L and Pelikan S 1985 Numerical computation of invariant circles of maps *Physica D* **16** 243–51
- [26] Hirsch M W 1994 *Differential Topology* (Berlin: Springer)
- [27] Hirsch M W, Pugh C C and Shub M 1977 *Invariant Manifolds* (Berlin: Springer)
- [28] Krauskopf B and Osinga H M 2003 Computing geodesic level sets on global (un)stable manifolds of vector fields *SIAM J. Appl. Dyn. Syst.* **4** 546–69
- [29] Krauskopf B, Osinga H M, Doedel E J, Henderson M E, Guckenheimer J M, Vladimirovsky A, Dellnitz M and Junge O 2005 A survey of methods for computing (un)stable manifolds of vector fields *Int. J. Bifurc. Chaos* **15** 763–91
- [30] Kuznetsov Y 1998 *Elements of Applied Bifurcation Theory* (Berlin: Springer)
- [31] Lang S 2002 *Introduction to Differentiable Manifolds* (Berlin: Springer)

- [32] Lorenz E 1979 *On the Prevalence of Aperiodicity in Simple Systems (Lecture Notes in Mathematics vol 755)* ed M Grmela and J E Marsden (Berlin: Springer) pp 53–75
- [33] Lorenz E 1984 Irregularity: a fundamental property of the atmosphere *Tellus* **36A** 98–110
- [34] Martin D 2002 *Manifold Theory: An Introduction for Mathematical Physicists* (Chichester: Ellis Horwood)
- [35] Maunder C 1970 *Algebraic Topology* (London: Van Nostrand-Reinhold)
- [36] Moore G 1996 Computation and parametrization of invariant curves and tori *SIAM J. Numer. Anal.* **33** 2333–58
- [37] Osinga H M 1996 Computing invariant manifolds *PhD Thesis* University of Groningen
- [38] Palis J and Takens F 1993 *Hyperbolicity & Sensitive Chaotic Dynamics at Homoclinic Bifurcations* (Cambridge: Cambridge University Press)
- [39] Phillips M, Levy S and Munzner T 1993 Geomview: an interactive geometry viewer *Not. Am. Math. Soc.* **40** 985–8
- [40] Press W, Teukolsky S, Vetterling W and Flannery B 1992 *Numerical Recipes in C* (Cambridge: Cambridge University Press)
- [41] Reichelt V 2000 Computing invariant tori and circles in dynamical systems *IMA Vol. Math. Appl.* **119** 407–37
- [42] Roussel M R and Fraser S J 1991 On the geometry of transient relaxation *J. Chem. Phys.* **94** 7106–13
- [43] Ruelle D 1989 *Elements of Differentiable Dynamics and Bifurcation Theory* (Boston, MA: Academic)
- [44] Sacker R 1969 A perturbation theorem for invariant manifolds and Hölder continuity *J. Math. Mech.* **18** 705–62
- [45] Schilder F, Osinga H M and Vogt W 2005 Continuation of quasi-periodic invariant tori *SIAM J. Appl. Dyn. Syst.* **4** 459–88
- [46] Schilder F, Osinga H M, Vogt W and Schreiber S 2006 Fourier methods for quasi-periodic oscillations *Int. J. Numer. Methods Eng.* **67** 629–71
- [47] Shil'nikov A, Nicolis G and Nicolis C 1995 Bifurcation and predictability analysis of a low-order atmospheric circulation model *Internat. J. Bifurc. Chaos Appl. Sci. Eng.* **5** 1701–11
- [48] Simó C 1998 Effective computations in celestial mechanics and astrodynamics *Modern Methods of Analytical Mechanics and their Applications* ed V V Rumyantsev and A V Karapetyan (Vienna: Springer) pp 55–102
- [49] Sparrow C 1982 *The Lorenz equations: bifurcations, chaos and strange attractors* (New York: Springer)
- [50] Strang G 1988 *Linear Algebra and its Applications* (San Diego, CA: Harcourt Brace Jovanovich)
- [51] Whitney H 1957 *Geometric Integration Theory* (Princeton, NJ: Princeton University Press)
- [52] Wong Y 1967 Differential geometry of Grassmann manifolds *Proc. Natl. Acad. Sci.* **57** 589–94